CONTROL DATA

Magnetic Storage Concepts

Version W

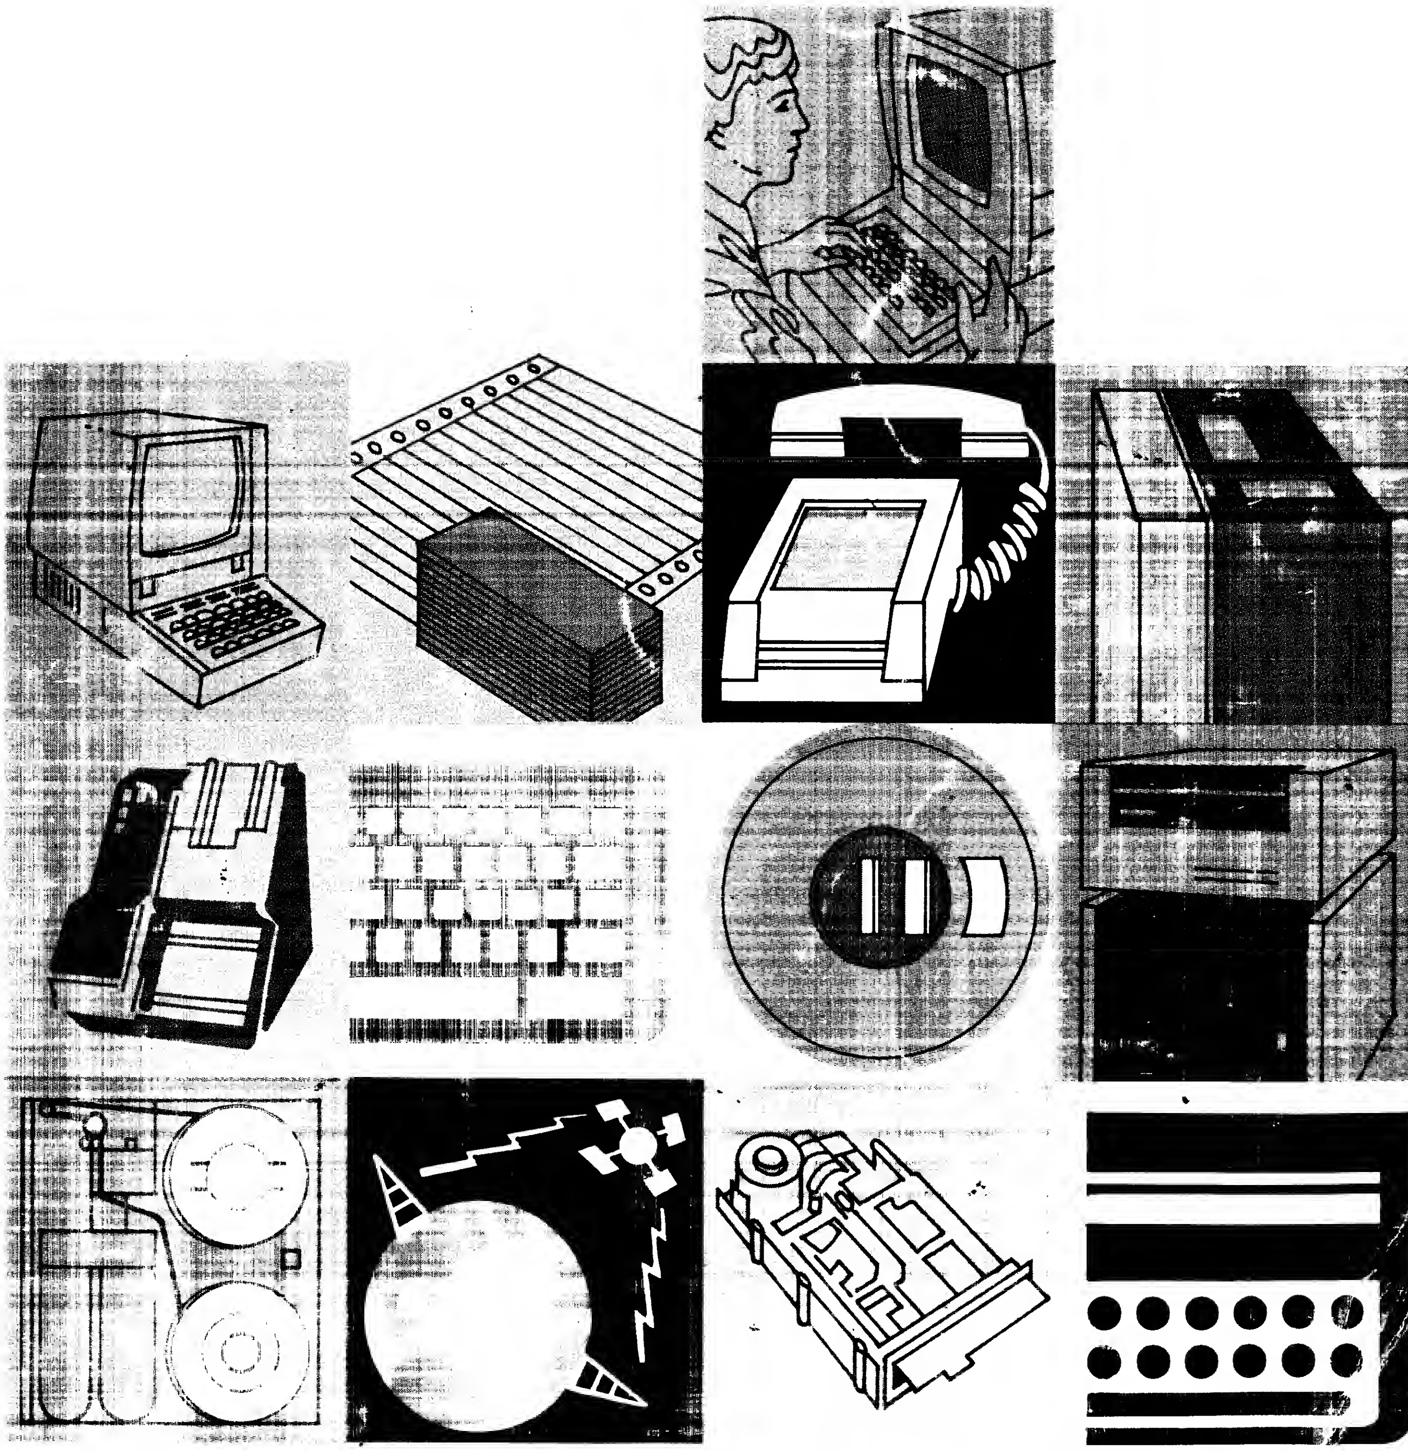

REFERENCE MANUAL

Magnetic Storage Concepts

CONTROL DATA

Pub. No. 76361245 A

Printed in the United States of America.
2/83

Table of Contents

Block 5
TAPE TRANSPORT MECHANICS

Block 6
TAPE TRANSPORT MECHANICS CONTROL

Block 7
TAPE SENSORS AND READ/WRITE HEADS

Block 8
INTRODUCTION TO DISK STORAGE DEVICES

Block 9
MAGNETIC RECORDING

Block 10
MOVABLE HEADS

Block 11
CONTROL LOGIC

Block 12
OTHER MAGNETIC STORAGE DEVICES

Block 1

Introduction to Magnetic Recording

Introduction to Magnetic Principles

Types of Magnets

Man's first experience with magnetism involved the mineral magnetite known to the ancient Greeks. The Greeks noticed that stones containing magnetite could attract pieces of iron. Magnetite, then, is regarded as a natural magnet. However, such natural magnets are too weak and inefficient for our present needs.

Today, most of the magnets that are used are created by artificial means. Often, they are produced with material that is capable of being magnetized, such as iron or steel. This process involves wrapping many turns of insulated wire around the material (see figure 1-1).

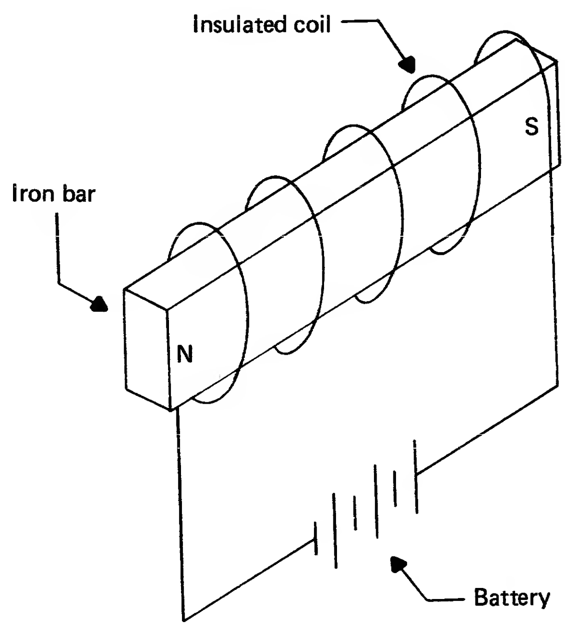

Figure 1-1. Forming an artificial magnet

When a direct current is passed through the wire, the material becomes magnetized. Certain types of material, such as hardened steel or certain alloys, will retain their magnetism when the current is removed and can become permanent magnets. Other materials lose their magnetism once the current is removed and are referred to as temporary magnets.

Principles of Magnetism

All magnets have a north pole and a south pole with a magnetic field between the two. Figure 1-2 shows a bar magnet with its field indicated by imaginary lines of force. These lines of force are believed to flow from the north pole to the south pole. When the lines of force are referred to collectively, they are called the magnetic flux. Flux density is the measurement of the number of lines of force in a given area.

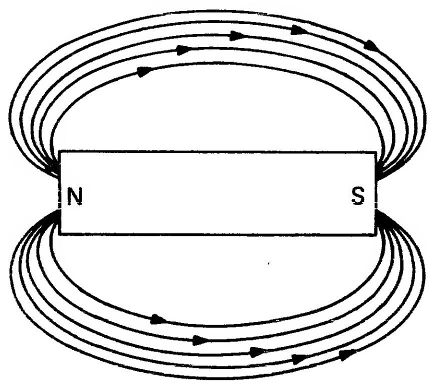

Figure 1-2. Magnetic lines of force

Looking at the iron bar at the molecular level, you can see that its molecules have a north/south polarity. In an unmagnetized piece of iron, the positions of these molecules are random, and, overall, their tiny fields cancel each other (see figure 1-3). In a magnetized bar, the molecules are all aligned in the same direction, and the effect of their fields is multiplied (see figure 1-4).

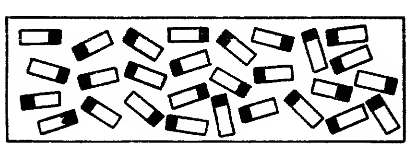

Figure 1-3. Molecular theory of magnetism, unmagnetized

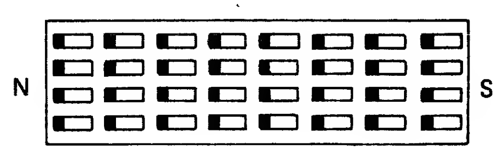

Figure 1-4. Molecular theory of magnetism, magnetized

Electromagnets

Whenever a current passes through a piece of wire, a small magnetic field builds up around the wire. As was mentioned earlier, wrapping several turns of wire around a steel bar produces a magnet when current passes through the wire. If the direction of the current is reversed, the polarity of the magnet will reverse. In other words, the end that was the north pole will become the south pole, and the south pole will become the north pole.

Electromagnets need not be bars; they have other physical shapes. Figure 1-5 shows a horseshoe magnet, which is similar to the type used in magnetic recorders to read and write information. In a horseshoe magnet, the poles are physically closer together than in a bar magnet. As a result, the field between the poles is relatively stronger than that in a bar magnet. This is because the strength of a magnetic field diminishes considerably as the distance between the poles increases.

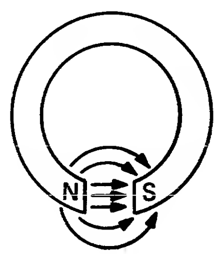

Figure 1-5. Horseshoe magnet showing concentration of magnetic field

There is another interesting property of this horseshoe type electromagnet. If another magnet is passed beneath the horseshoe magnet, an electrical current will be produced in the wire of the coil (see figure 1-6). This property of the generation of electrical current by the interaction of different magnetic fields is known as induction. Induction is the principle that is in effect when read heads are used in magnetic surface recording.

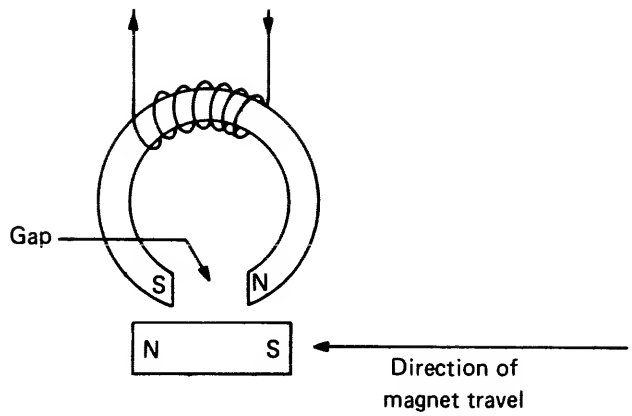

Figure 1-6. Inducing current in a coil

Magnetic Surface Recording

These magnetic principles can now be applied to surface recording. In surface recording, a surface of substrate is coated with a thin coating of material capable of being magnetized and retaining that magnetism indefinitely. You will now look at a hypothetical write head.

In figure 1-7, you see that the write head is essentially an electromagnet in the shape of a horseshoe; there is no electrical current in the coil, and no magnetism.

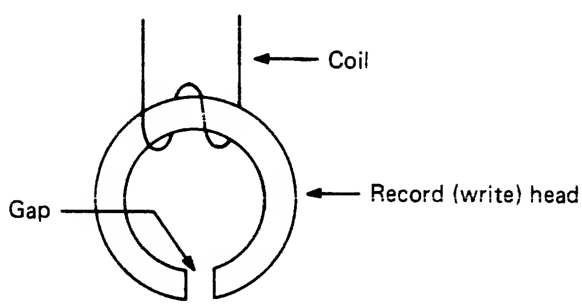

Figure 1-7. Recording head, no current through coil (simplified drawing)

In figures 1-8 and 1-9, you see the head with current flow. The poles of the magnet produce a magnetic field that is strong enough to cause changes in materials close to the head. If a magnetic surface is passed beneath the head while current is flowing in the coil of the head, the particles in the surface will be aligned with the same polarity as the head.

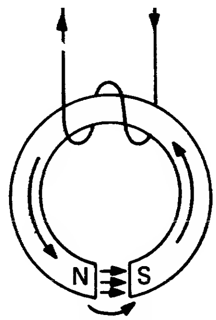

Figure 1-8. Recording head, current through coil (simplified drawing)

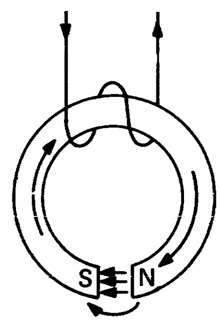

Figure 1-9. Recording head, current reversed (simplified drawing)

Now, if the direction of the current flow is changed, the polarity of the head will be changed, as will the polarity of the material passing beneath the head (see figures 1-10, 1-11, and 1-12). Since the effect of the magnetic field decreases rapidly as the material is moved away from the magnet, the effect on the recording surface is confined to the immediate area of the head.

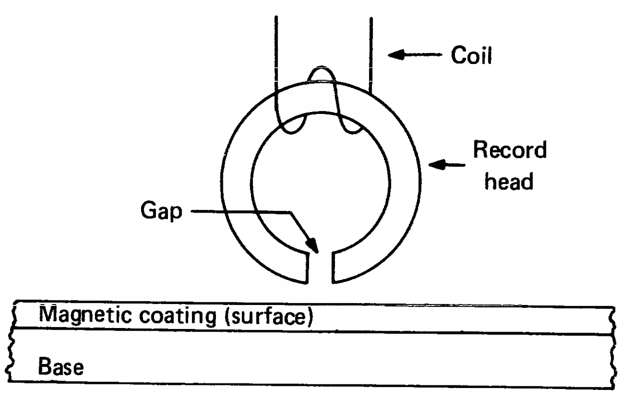

Figure 1-10. Recording on a magnetic surface, no current through coil

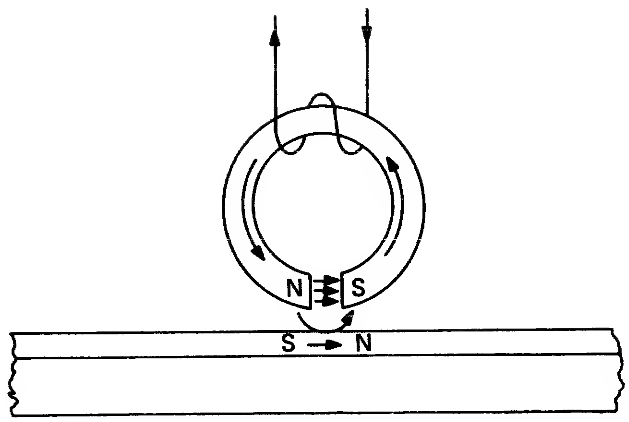

Figure 1-11. Recording on a magnetic surface, current through coil

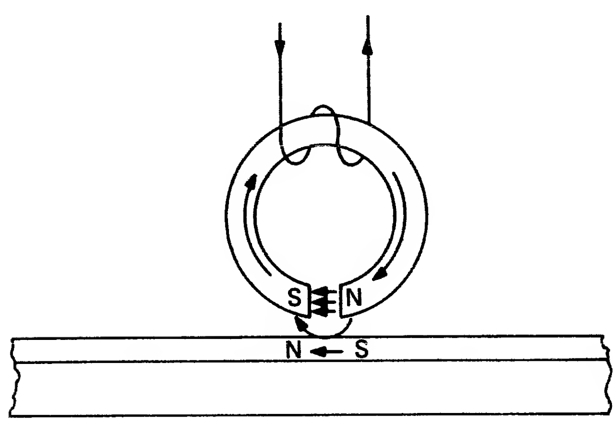

Figure 1-12. Recording on a magnetic surface, current reversed

Reading the Magnetic Surface

These polarity changes recorded on the magnetic surface can be detected with a read head. The read head is very similar in construction to the write head. As a previously recorded surface passes beneath the read head, any changes in magnetic polarity on the surface cause a current to flow in the coil of the read head (see figure 1-13). The direction of the polarity change on the surface determines the direction of current flow in the wire. It should be noted that if there is no change in polarity on the surface, no current will be induced in the coil. Also, the current produced by induction is very small when compared with the current required for a write operation.

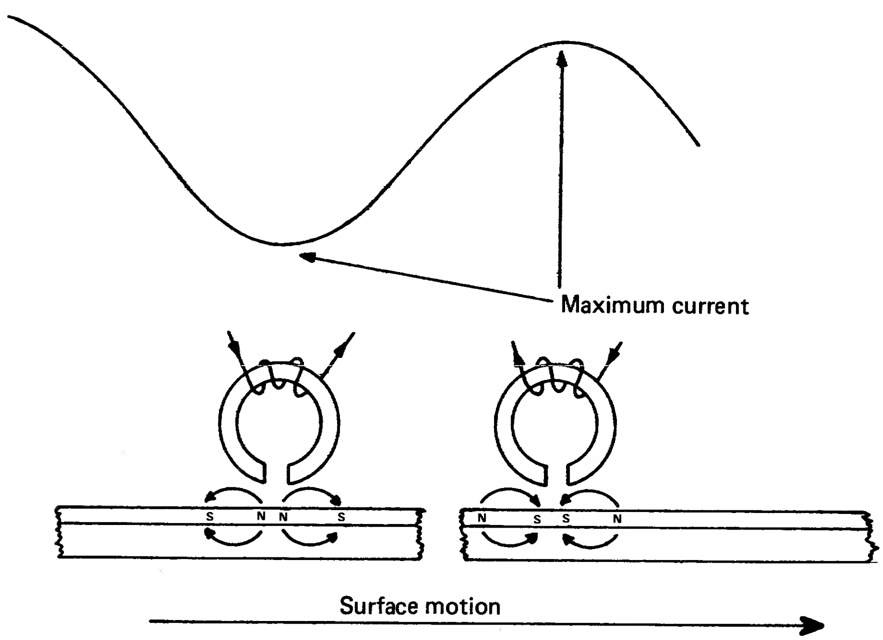

Figure 1-13. Relative head to surface motion, reproducing (read operation)

Reading the magnetic surface does not alter the changes on the surface. As a result, the same area can be reread indefinitely.

Tape Construction (Text)

Magnetic Tape Components

The magnetic tape used in digital recording is composed of three basic parts: 1) the oxide coating, 2) the binder, and 3) the base, or backing material.

Oxide Coating

The oxide coating used is composed of gamma-ferric oxide, Fe_2O_3, in the form of tiny, needle-shaped particles less than 1 micron in length. A micron is .000039, or 39/1,000,000, of an inch. The oxide forms a thin (.00045 inch) layer on the tape, and the layer must be uniform for accurate recording. A tape must have an oxide coating that varies no more than .000015 inch in thickness over its entire length.

In addition, the oxide coating used in digital recording must have the following characteristics:

- It must be capable of being easily magnetized.
- It must retain its magnetism indefinitely so that the tapes can be stored for long periods of time.
- It must show a high resistance to stray magnetic fields. (This prevents accidental erasures and allows the tiny areas of recorded data to be packed closely together.)

Binder

This oxide coating is bound to the tape with the binder, which normally is made up of lubricants, dispersing agents, and solvents. The composition of the binder is very important in the manufacture of the tape. To prevent the oxide coating from chipping and flaking, the binder must be both tough and flexible. Therefore, the individual formulas for the binder are protected by patent by each tape manufacturer.

Backing

The backing commonly used in tape manufacture is a polyester film called Mylar, which is a DuPont trade name for polyethylene terephthalate. Its main advantages are its strength, humidity stability, and resistance to solvents. This polyester film base is only .00142 inch, or 1.42 mils, thick.

Tape Dimensions

When tape that is very thin is wound on a reel, a condition called "print through" can occur. In such an instance, the magnetized areas of oxide on one layer may exert their magnetism on the adjacent layer and cause flux changes in that layer. Thus, the tape would have data errors.

On the other hand, if the tape is excessively thick, much less tape will fit on a reel. To handle large reels of this heavier tape, larger and more expensive motors would be required in the reel drive and tape drive systems of the tape transports. Because of the sudden starts and stops performed by high-speed tape transports, the tape used must be as light as possible to reduce excessive inertia when the tape is moved or stopped. These sudden starts and stops also impose the requirement that the tape be strong, to withstand stretching or breaking.

Digital magnetic tapes are cut to a width of one-half inch, and come in varying lengths. The most common lengths are 300-, 800-, 1200-, and 2400-foot reels. Of these, the 2400-foot length is the most commonly used, and is wound on a plastic reel with an outside diameter of 10.5 inches.

Certification

After a tape is manufactured, its entire length is checked for any spots that would prevent the accurate recording of data. This process is known as "certification." Since recording frequencies vary according to the machine used, the certification for each tape is done at the frequency for which the tape is intended. Each tape then receives a label with phrases such as "Certified 1600 Characters Per Inch (CPI)," "Certified 3200 Frames Per Inch (FRPI)," or "Certified 800 Bits Per Inch (BPI)," thus indicating the maximum frequency at which the tape should be used.

BOT/EOT Markers

Each digital magnetic tape also has two reflective markers on it to indicate the beginning and end of the tape. These markers are about an inch long and two-tenths of an inch wide and are shiny on the outside. The beginning-of-tape, or BOT, marker is placed about 16 feet from the beginning of the tape on the side that does not have oxide. It is sometimes referred to as the load point. It is placed on the edge of the tape that is normally toward the front portion of the reel. The end-of-tape, or EOT, marker is similar, but is placed 25 feet from the end of the tape and on the edge of the tape toward the back of the reel. These markers enable the tape transport to determine the points at which it should start and stop recording.

Block 2

Recording Methods and Codes

Recording Methods (Text)

Although magnetic recording techniques are used in a number of areas, at this time, you will only be involved with the recording and retrieval of digital information–that is, information represented by binary 1s and 0s. There are many techniques used in recording digital information, but you are going to look only at the following four methods commonly used in the data processing industry.

- NRZI – Nonreturn to Zero, Indiscrete method. This was the first widely used technique for digital tape recording in the computer industry, and is still in use today.
- PE – Phase Encoding method (also known as Phase Modulation method). This method allows data to be recorded at higher densities than are permitted by the standard NRZI format.
- GCR – Group Coded Recording Method. This is not actually a separate recording technique, but is a highly sophisticated method of formatting data to allow NRZI recording at extreme densities.

  (The preceding three methods are found most often in digital magnetic tape systems.)

- FM – Frequency Modulation method. This technique is commonly found in recording magnetic disks, drums, and files.

NRZI Recording Method

In the discussion of NRZI recording, magnetic tape will be used as the media. Keep in mind, however, that the technique may be applied to any magnetic recording surface.

You probably recall from a previous activity that whenever a flux change occurs on tape passing beneath a read head, a current is produced in the coil of the read head. The direction of the current, either positive or negative, depends on whether the flux change is from a south/north to north/south configuration, or from north/south to south/north.

In NRZI recording, a binary 1 is indicated by a flux change, and a 0 by no flux change. In this type of recording, the direction of the flux change does not matter; it is important only that a flux change occurred (see figure 2-1).

In writing NRZI data, it is necessary to change the direction of the current in the write head coil in order to record a 1. To record a 0, the direction of the current is not changed, but is maintained in the same direction it was after the last change.

To read data recorded in this format, the read logic will generate a 1 whenever current is induced in the read head coil; this indicates that a flux change was detected. If there is no flux change detected in the read coil, a 0 is generated.

Consider the problem of reading back a string of recorded 0s. Since there are no flux changes, there is no way to tell how many 0s are actually recorded in a string of 0s. To avoid this problem, the read logic must generate timing pulses to determine when it is reading data; this sometimes is called a "read window." The read logic will detect a 1 or a 0 only at the time of a read window. This need for an external clock scheme in reading NRZI data is considered one of the method's disadvantages.

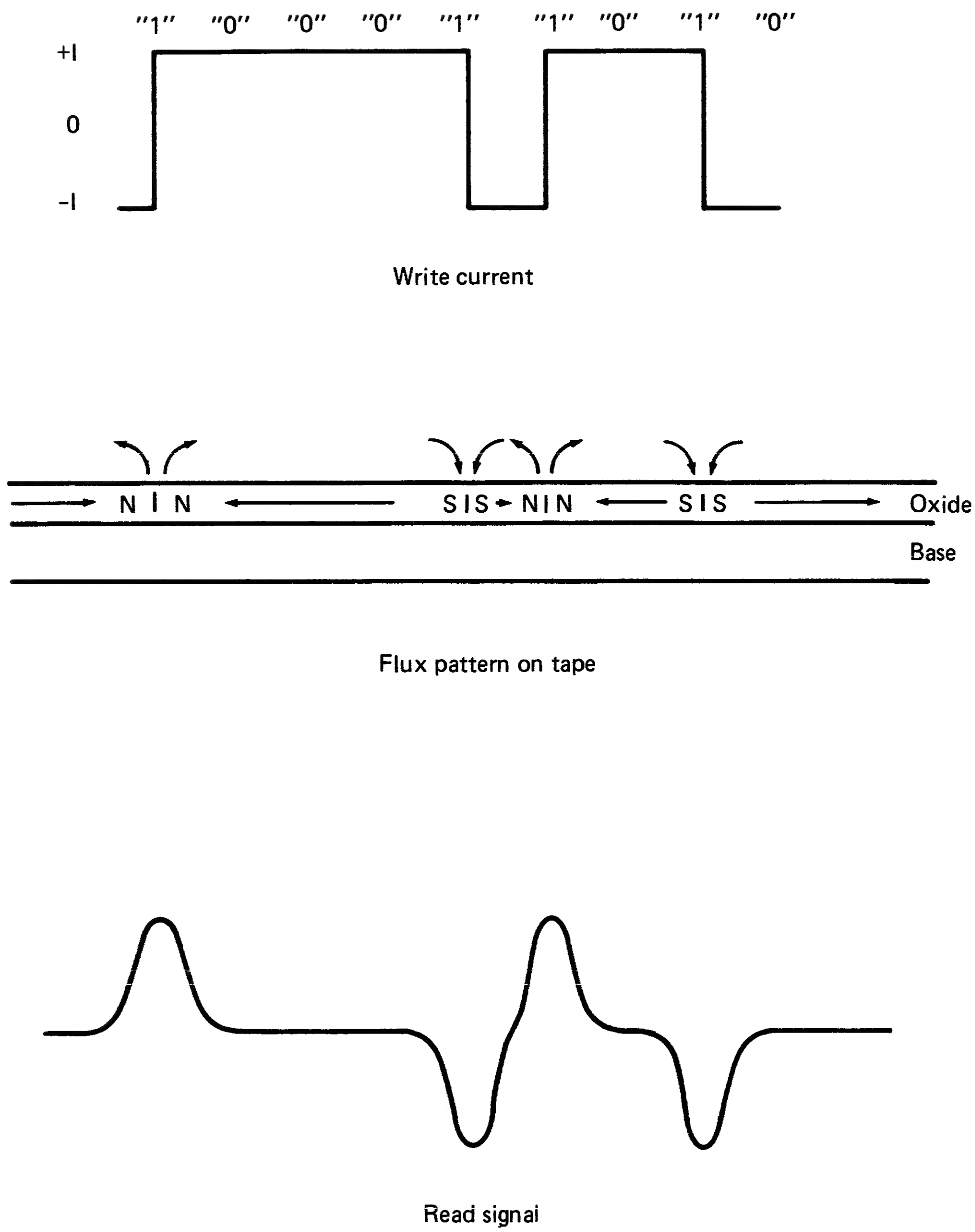

Figure 2-1. NRZI recording

Phase Encoding Method

The phase encoding or phase modulation method of recording is more complicated than the NRZI method. In phase encoding, the direction of the current flow caused by the flux change determines whether the flux change represents a 1 or a 0. A positive current flow represents a 1, and a negative current flow represents a 0.

Decoding Phase Encoding Data

In order to decode phase encoding data, it must be broken down into cells. One cell is the time it takes to transfer one data bit. Flux changes in the center of the cell are considered the data. Figure 2-2 shows a phase waveform. The dashed vertical lines represent the edges of the cells. You will notice that whenever a positive-going pulse appears in the center of a cell, a 1 appears in that cell; whenever a negative-going pulse occurs, a 0 appears.

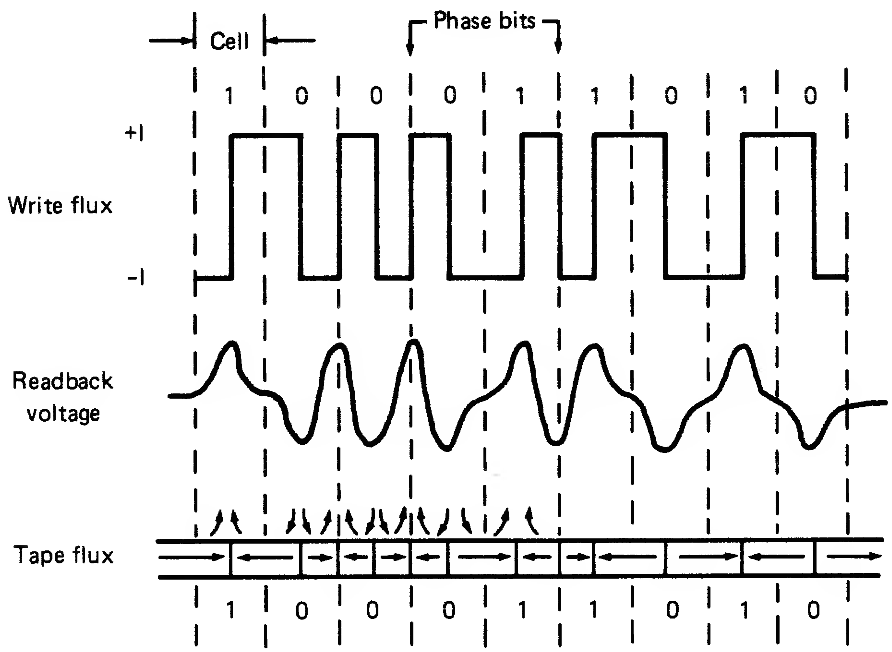

Figure 2-2. Phase encoding recording

If the data has two 0s in a row, you will need to induce two negative pulses in a row. To produce two negative pulses in a row, two of the same flux changes are needed. However, you cannot have an N/S to S/N change followed by another N/S to S/N change unless the flux is changed back to N/S in between. This is precisely what is done at the border of the two cells. In terms of write current, if you have to go from positive to negative in two consecutive cells, then between those cells you have to go back to positive. Recording two 1s in a row is done in the same way, except the flux reversals and current flows are the opposite. These flux reversals between cells are called "phase bits." If two cells in a row have different values, there is no need to reverse polarity between the cells, and no phase bits will appear.

Although these phase bits are not data bits, they can be used for error checking. For example, if you have just read a 1 in a cell and encounter a phase bit, you know that the next cell should contain a 1; likewise, if you have just read a 1 and encounter no phase bit, you know the next cell should contain a 0.

Since phase encoding requires either one or two flux changes for every data bit, recording is done at a very narrow range of frequencies. This makes it relatively easy to design a read/write head that is tuned to operate at this frequency range. The frequency of NRZI recording depends on the actual data and will vary greatly, thus making it more difficult to design a tuned read/write head. Because of this difference, heads for phase encoding can be designed to operate at frequencies higher than those at which NRZI heads operate. Consequently, phase encoding can be performed at a higher frequency than NRZI recording, producing more data (bit density) in the same area.

These advantages—higher densities and error correction potential—are achieved at the expense of more complicated read logic.

Group Coded Recording Method

This activity introduces the group coded recording (GCR) method of recording by first discussing the advantages and disadvantages of phase encoding and NRZI. By combining the advantages of phase encoding (PE) and NRZI, this new method, GCR, resulted.

PE and NRZI Methods

The phase encoding (PE) method and the NRZI method of recording data on tape each has its own advantages and disadvantages. The advantage of the phase encoding method is that it is self-clocking; that is, there is always a signal present on the tape, even if the data happens to be a string of 0s.

This signal makes it much easier for the hardware to detect the correct information on the tape while rejecting unwanted noise. The disadvantage of this method is that the recording frequency is twice as high as the frequency of the data on tape. Thus, a high-quality tape, capable of recording 3200 bits per inch, would only be able to have 1600 bits per inch of data. The NRZI method is more efficient because it records only 1 bits on tape, leaving blank spaces for 0 bits.

The disadvantage of the NRZI method, of course, is that the individual tracks cannot be self-clocking, since a string of successive 0s produces no signal.

GCR Method

The group coded recording (GCR) method combines the advantages of the above systems to produce extremely high-density recordings. In GCR, the recording method is actually the same as NRZI. However, instead of writing the data on the tape in its true form, the data is first sent through a coder. (See figure 2-3.)

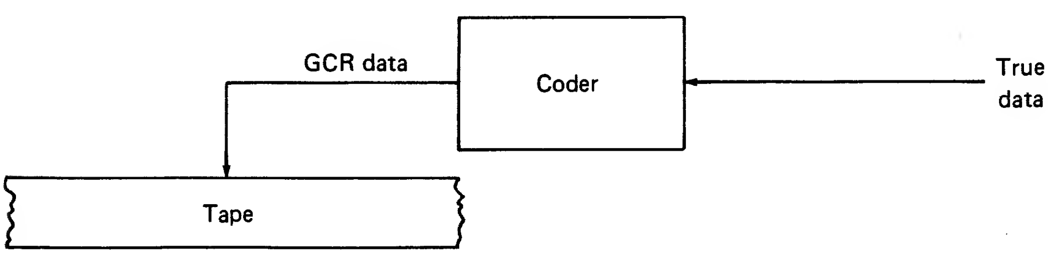

Figure 2-3. GCR recording

The coder converts the true data into a series of 5 bit codes. These codes will never have more than two consecutive 0s, regardless of what the true data was. This is important because if there are never more than two consecutive 0s, the data can be self-clocking.

When the GCR data is read from the tape, it is sent through the coder again and converted back to its true form. (See figure 2-4.)

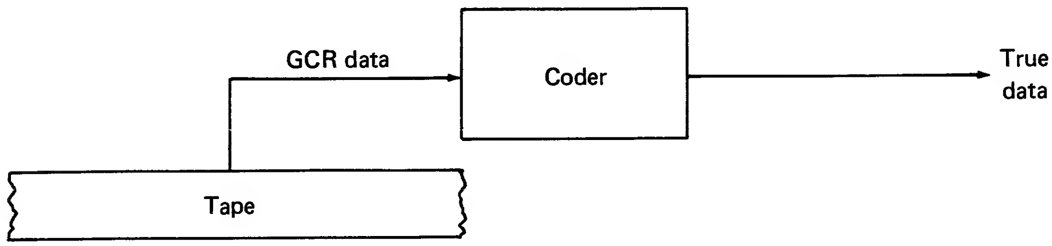

Figure 2-4. GCR data decoding

Using this method, it is possible to achieve recording densities in excess of 6000 bits per inch.

Table 2-1 shows some typical data values along with their equivalent coded values. In order to have a code that has no more than two consecutive 0s, it is necessary to add an extra bit to the code. Therefore, every 4 bits of true data will have a 5-bit GCR code.

TABLE 2-1
Record Code Values

Data Values	Record Values
Group Position: 1234/5678	Group Position: 12345/678910
0000	11001
0001	11011
0010	10010
0011	10011
0100	11101
0101	10101
0110	10110
0111	10111
1000	11010
1001	01001
1010	01010
1011	01011
1100	11110
1101	01101
1110	01110
1111	01111

In addition to having extremely high record densities, the GCR method is also capable of accurately detecting and correcting a wide range of errors. In most cases, the error correction can be done "on the fly," meaning the tape does not have to be stopped and reread when an error is encountered. Features such as these require very sophisticated logic design and precision transport mechanics.

A special format for the recorded information is also required (see figure 2-5). The following terms appear in figure 2-5.

ID Burst – A series of bits in track 6 that tell the logic circuits that the tape has been recorded in GCR mode.

ARA Burst – A series of bits on all tracks that automatically adjust the gain of the read amplifiers to match the signal level of the tape.

ID Character – A special code that tells the logic to get ready to read data.

Data Blocks – GCR data.

Tape Mark – A special code that identifies the end of a group of blocks.

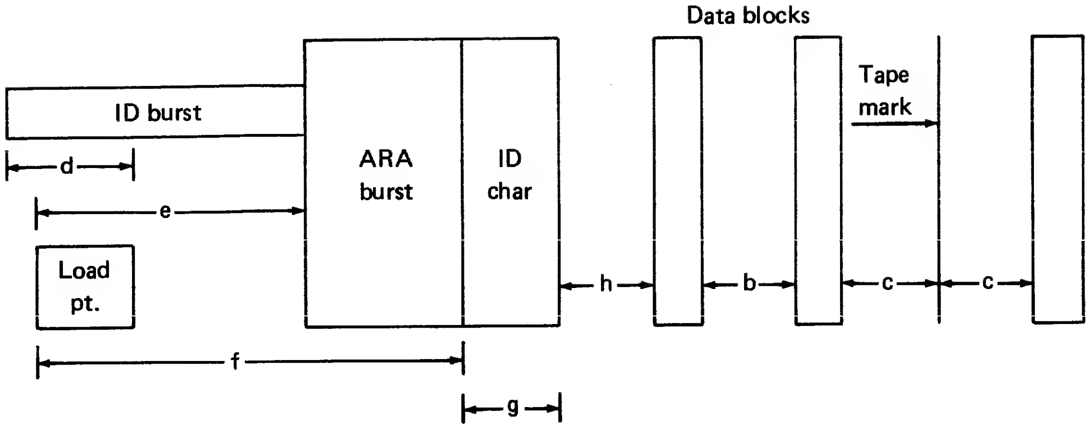

Inter-block gap (b)		.25 in. minimum .30 in. nominal .35 in. maximum
Filemark gap (c)		3.0 in.
Dimension	d	2.3 in.
	e	2.0 in.
	f	10.5 in.
	g	2.0 in.
	h	2.5 in.

Tape marks consist of 320 ± 10% flux reversals at 9042 fci in tracks 2, 5, 8, 1, 4, and 7. Tracks 3, 6, and 9 are DC erased.

The ID burst is written at 3014 fci.

Figure 2-5. GCR recording specifications

Frequency Modulation Method

Figure 2-6 shows the waveform for a signal produced by the frequency modulation recording technique. As you can see, a flux reversal at the center of a cell indicates a 1, while no flux reversal at the center of the cell indicates a 0. There will always be a flux reversal between the cells. As in NRZI recording, you are concerned not with the direction of the flux reversals, but with whether a flux change has occurred.

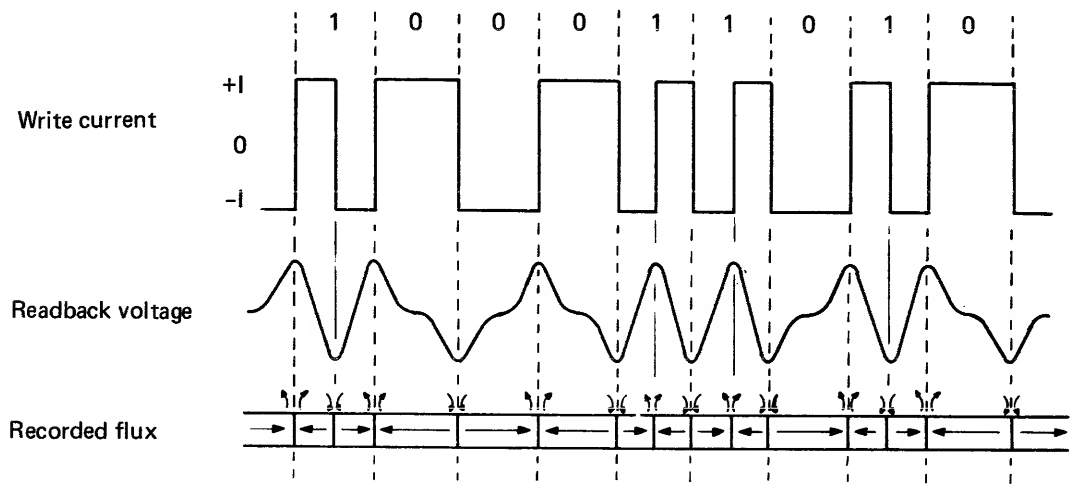

Figure 2-6. Frequency modulation

As with phase encoding, there will always be one or two flux changes per bit. Thus, read/write heads that are tuned to operate at higher frequencies than NRZI heads are produced, allowing FM recording to be performed at higher frequencies.

Since there will always be a flux reversal between cells, this method is self-clocking and the read circuits are simplified.

Coding Methods

Basically, all digital computers are binary machines; that is, operations within the machines, as well as data from the machines, are based on 1s and 0s. But while binary data holds great significance for the computer, a string of 1s and 0s means little to a human being.

Therefore, coding systems were developed to translate this binary information into data that is usable by people. In all of these coding systems, specific binary data was selected to symbolize specific characters. For example, given the binary information 111, you could assign this data to the numeral 7. Then, whenever the computer gave the data 111, you could convert it to 7. If you wanted to represent a letter, you could do so easily. Given the code for the letter A, you could assume that whenever the computer gave the binary data 010001, it would represent the character A.

This is essentially what computer manufacturers did as they developed their machines. Unfortunately, manufacturers frequently developed their codes based on characteristics particular to their own machines. As a result, the exchange of data among machines made by different manufacturers was greatly hampered by different coding systems. Attempts were made over the years to standardize coding systems, but even today there are several different codes in use.

Because all coding systems are based on binary data, the number of different characters that can be represented in any system is limited by the number of binary bits used to represent a character. With a 3-bit code, you could represent only 8 characters since only eight combinations are possible with 3 binary bits. Likewise, with a 6-bit code, you would have 64 possible characters; with a 7-bit code, 128 characters; with an 8-bit code, 256 characters.

You will examine three different coding systems: binary coded decimal, ASCII, and EBCDIC. These systems, though widely used in the data processing industry, represent only a small fraction of the coding systems in existence today.

BCD Code

BCD stands for Binary Coded Decimal. The most common BCD codes are 6-bit codes, although there is a 4-bit BCD code that is used for decimal arithmetic in computer operations to represent the values 0 to 9.

A 6-bit BCD code allows for the possibility of only 64 different characters; thus, it eliminates the possibility of including both uppercase and lowercase letters. Normally BCD includes only uppercase letters, numbers, and a few special characters. While that is sufficient for many applications, more and more cases are requiring a code with a greater number of characters.

Table 2-2 contains BCD codes. You will notice columns labeled External BCD and Internal BCD. External BCD is the common BCD coding that is recorded. However, external BCD codes do not lend themselves to convenient sorting. For example, an external BCD A has a value of 61, which is higher than a Z with a value of 31. On the other hand, the letter S has a value of 22, which is lower than a Z. To avoid this problem, external BCD is usually converted to internal BCD when it is being read into the computer. In internal BCD, the upper bit of certain characters is complemented so that the alphabetic characters are in ascending numerical order (see table 2-3). When data is being recorded, the internal BCD can be converted back to external BCD.

Looking at the table of BCD characters in table 2-2, you can see that the character G is represented by the internal BCD code of 27. Given a BCD code of 44, you should be able to find the character that it represents (the letter M).

TABLE 2-2
Magnetic Tape BCD Codes

Internal BCD Codes	External BCD Codes	Magnetic Tape BCD Codes	Hollerith Card Codes	Key Punch Characters	Printer Characters
00	12	82	0	0 (zero)	0
01	01	C1	1	1	1
02	02	C2	2	2	2
03	03	21	3	3	3
04	04	C4	4	4	4
05	05	41	5	5	5
06	06	42	6	6	6
07	07	C421	7	7	7
10	10	C8	8	8	8
11	11	81	9	9	9
12	(illegal)		8, 2	---	---
13	13	C821	8, 3	#	=
14	14	84	8, 4	@	≠
15	15	C841	8, 5	---	≤
16	16	C842	8, 6	---	%
17	17	8421	8, 7	(tape mark)	[
20	60	BA	12	+	+
21	61	CBA1	12, 1	A	A
22	62	CBA2	12, 2	B	B
23	63	BA21	12, 3	C	C
24	64	CBA4	12, 4	D	D
25	65	BA41	12, 5	E	E
26	66	BA42	12, 6	F	F
27	67	CBA421	12, 7	G	G
30	70	CBA8	12, 8	H	H
31	71	BA81	12, 9	I	I
32	72	BA82	12, 0	+0	<
33	73	CBA821	12, 8, 3	.	.
34	74	BA84	12, 8, 4	)	)
35	75	CBA841	12, 8, 5	---	≥
36	76	CBA842	12, 8, 6	---	?
37	77	BA8421	12, 8, 7	---	;
40	40	CB	11	- (minus)	-
41	41	B1	11, 1	J	J
42	42	B2	11, 2	K	K
43	43	CB21	11, 3	L	L
44	44	B4	11, 4	M	M
45	45	CB41	11, 5	N	N
46	46	CB42	11, 6	O	O
47	47	B421	11, 7	P	P
50	50	B8	11, 8	Q	Q
51	51	CB81	11, 9	R	R
52	52	CB82	11, 0	-0	V
53	53	B821	11, 8, 3	S	S
54	54	CB84	11, 8, 4	*	*

TABLE 2-2–Continued

Internal BCD Codes	External BCD Codes	Magnetic Tape BCD Codes	Hollerith Card Codes	Key Punch Characters	Printer Characters
55	55	B841	11, 8, 5	---	)
56	56	B842	11, 8, 6	---	;
57	57	CB8421	11, 8, 7	---	∧
60	20	CA	(blank)	(space)	
61	21	A1	0, 1	/	/
62	22	A2	0, 2	S	S
63	23	CA21	0, 3	T	T
64	24	A4	0, 4	U	U
65	25	CA41	0, 5	V	V
66	26	CA42	0, 6	W	W
67	27	A421	0, 7	X	X
70	30	A8	0, 8	Y	Y
71	31	CA81	0, 9	Z	Z
72	32	CA82	0, 8, 2	(record mark)	]
73	33	A821	0, 8, 3	(comma)	,
74	34	CA84	0, 8, 4	(	(
75	35	A841	0, 8, 5	---	→
76	36	A842	0, 8, 6	---	≡
77	37	CA8421	0, 8, 7	---	∧

TABLE 2-3
Internal-External BCD Codes

Char Prd.	Int. BCD	Ext. BCD
A	21	61
B	22	62
C	23	63
D	24	64
E	25	65
F	26	66
G	27	67
H	30	70
I	31	71
J	41	41
K	42	42
L	43	43
M	44	44
N	45	45
O	46	46
P	47	47
Q	50	50
R	51	51
S	62	22
T	63	23
U	64	24
V	65	25
W	66	26
X	67	27
Y	70	30
Z	71	31

ASCII Code

American Standards Code Information Interchange, or ASCII, often pronounced "as-key," was introduced in 1963 by the American Standards Association (now the American National Standards Institute) as a 7-bit standard code intended to provide a standardized communications medium. This code is sometimes referred to as ANSI; however, the title ANSI is incorrect.

The 7-bit code provides a possibility of 128 different characters—twice as many as provided by the 6-bit BCD.

Table 2-4 shows the ASCII characters. If you want to find out what bit pattern represents the uppercase G, you look for a capital G on the chart. Next to it, under the ASCII bit representation, you see that it is represented by the binary value 100 0111. If, on the other hand, you are given the binary value 101 0111 and wish to see what ASCII character that represents, you first locate the binary value in the bit representations, and, following across the line, see that it stands for the capital W.

TABLE 2-4
ASCII Codes

CHARACTER	HOLLERITH	BCD	INTERNAL BCD	EBCDIC	ASCII
0	0	001010	000000	11110000	0110000
1	1	000001	000001	11110001	0110001
2	2	000010	000010	11110010	0110010
3	3	000011	000011	11110011	0110011
4	4	000100	000100	11110100	0110100
5	5	000101	000101	11110101	0110101
6	6	000110	000110	11110110	0110110
7	7	000111	000111	11110111	0110111
8	8	001000	001000	11111000	0111000
9	9	001001	001001	11111001	0111001
A	12-1	110001	010001	11000001	1000001
B	12-2	110010	010010	11000010	1000010
C	12-3	110011	010011	11000011	1000011
D	12-4	110100	010100	11000100	1000100
E	12-5	110101	010101	11000101	1000101
F	12-6	110110	010110	11000110	1000110
G	12-7	110111	010111	11000111	1000111
H	12-8	111000	011000	11001000	1001000
I	12-9	111001	011001	11001001	1001001
J	11-1	100001	100001	11010001	1001010
K	11-2	100010	100010	11010010	1001011
L	11-3	100011	100011	11010011	1001100
M	11-4	100100	100100	11010100	1001101
N	11-5	100101	100101	11010101	1001110
O	11-6	100110	100110	11010110	1001111
P	11-7	100111	100111	11010111	1010000
Q	11-8	101000	101000	11011000	1010001
R	11-9	101001	101001	11011001	1010010
S	0-2	010010	110010	11100010	1010011
T	0-3	010011	110011	11100011	1010100
U	0-4	010100	110100	11100100	1010101
V	0-5	010101	110101	11100101	1010110
W	0-6	010110	110110	11100110	1010111
X	0-7	010111	110111	11100111	1011000
Y	0-8	011000	111000	11101000	1011001
Z	0-9	011001	111001	11101001	1011010

EBCDIC Code

EBCDIC, frequently pronounced "ib'-si-dick," is the acronym for Extended Binary Coded Decimal Interchange Code. It is an 8-bit code that was developed by IBM for its 360/370 computer models.

This 8-bit code allows the possibility of 256 codes. This quantity of codes permits the coding of uppercase and lowercase alphabetics, numerics, and many special characters. In addition, the first 40 characters are special control characters used in data communications work.

The 8 bits of the code can be divided into two groups of 4 bits each. The left-hand group represents zone bits while the right-hand group represents digit bits. The zone bits are similar to the zone punches in punched cards; for example, a zone of 1100 would correspond to a 12 punch, 1101 would be an 11 punch, 1110 a 0 punch, and 1111 a no-zone punch. The digit punches correspond to the normal digit punches on a punched card.

EBCDIC's 8-bit code can be represented by two hexadecimal digits. Hexadecimal refers to a numbering system based on 16, and each digit in the hexadecimal system can be represented by 4 bits. This code is convenient for IBM since its computers are designed as hexadecimal machines.

Table 2-5 shows EBCDIC codes. If you look for the EBCDIC code for the capital letter A, you see that it is represented by the binary code 1100 0001 and by the hexadecimal code C1. Take another example: suppose you have the binary code 1110 1001. If you locate that code in the table, you find that it corresponds to the capital letter Z.

TABLE 2-5
EBCDIC Codes

CHARACTER	HOLLERITH	BCD	INTERNAL BCD	EBCDIC	ASCII
0	0	001010	000000	11110000	0110000
1	1	000001	000001	11110001	0110001
2	2	000010	000010	11110010	0110010
3	3	000011	000011	11110011	0110011
4	4	000100	000100	11110100	0110100
5	5	000101	000101	11110101	0110101
6	6	000110	000110	11110110	0110110
7	7	000111	000111	11110111	0110111
8	8	001000	001000	11111000	0111000
9	9	001001	001001	11111001	0111001
A	12-1	110001	010001	11000001	1000001
B	12-2	110010	010010	11000010	1000010
C	12-3	110011	010011	11000011	1000011
D	12-4	110100	010100	11000100	1000100
E	12-5	110101	010101	11000101	1000101
F	12-6	110110	010110	11000110	1000110
G	12-7	110111	010111	11000111	1000111
H	12-8	111000	011000	11001000	1001000
I	12-9	111001	011001	11001001	1001001
J	11-1	100001	100001	11010001	1001010
K	11-2	100010	100010	11010010	1001011
L	11-3	100011	100011	11010011	1001100
M	11-4	100100	100100	11010100	1001101
N	11-5	100101	100101	11010101	1001110
O	11-6	100110	100110	11010110	1001111
P	11-7	100111	100111	11010111	1010000
Q	11-8	101000	101000	11011000	1010001
R	11-9	101001	101001	11011001	1010010
S	0-2	010010	110010	11100010	1010011
T	0-3	010011	110011	11100011	1010100
U	0-4	010100	110100	11100100	1010101
V	0-5	010101	110101	11100101	1010110
W	0-6	010110	110110	11100110	1010111
X	0-7	010111	110111	11100111	1011000
Y	0-8	011000	111000	11101000	1011001
Z	0-9	011001	111001	11101001	1011010

Block 3

Tape Format, Error Detection, and Tape Handling

Tape Format/Magna-See®

Data is recorded on magnetic tape in small flux reversals that cannot be seen by simply looking at the tape. However, these flux reversals can be made visible by using a liquid tape developer such as Magna-See. (See figure 3-1.)

Figure 3-1. Magna-See developer

Magna-See consists of a liquid that holds a fine iron powder in suspension. The iron particles are not dissolved in the liquid, so if the can has been sitting still, the particles will tend to collect on the bottom. Shake the can vigorously to put the particles back in suspension.

The tape is developed by dipping it in the liquid. Use a back and forth, see-saw motion. (See figure 3-2.)

Figure 3-2. Tape being dipped

When the portion of the tape you want developed is coated, set the tape aside and let the liquid evaporate. The iron powder will remain, and it will collect around the magnetic flux reversals, making them visible. The tape in figure 3-3 happens to be a 7-track NRZI format tape. The vertical marks represent 1 bits.

Figure 3-3. Developed tape

The iron particles will not stick to the tape and you can brush them away. But you can also make a permanent, visual record of the data recorded on the tape. This permanent record can be placed in a log book or used to show a customer how data is recorded.

To make this permanent record, place a piece of ordinary, transparent adhesive tape over the portion of the tape you have developed. (See figure 3-4.)

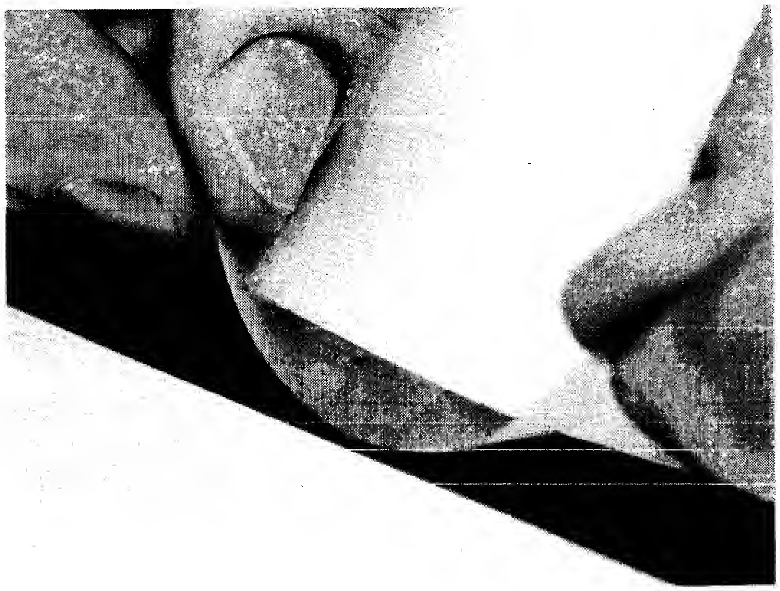

Figure 3-4. Transparent adhesive tape being applied

Press the tape firmly and then carefully peel it off. (See figure 3-5.)

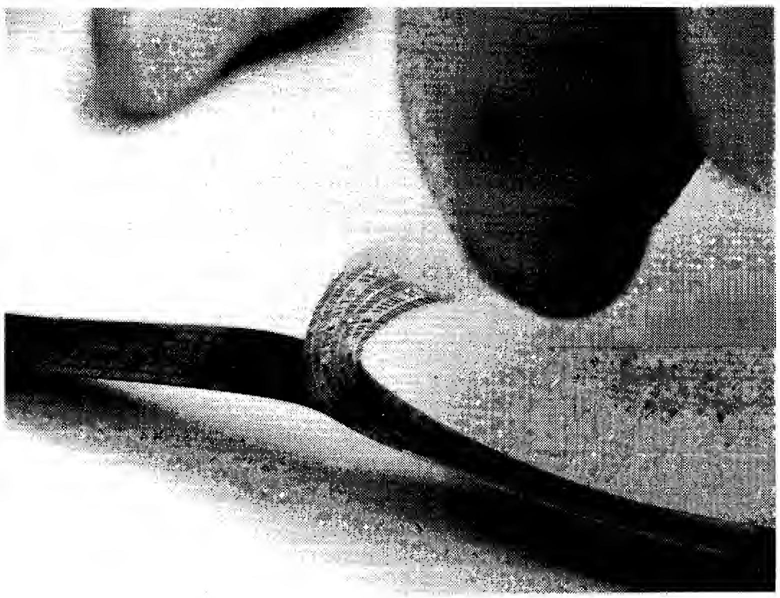

Figure 3-5. Transparent adhesive tape being removed

The iron powder will stick to the adhesive tape, leaving an accurate image of the data. (See figure 3-6.)

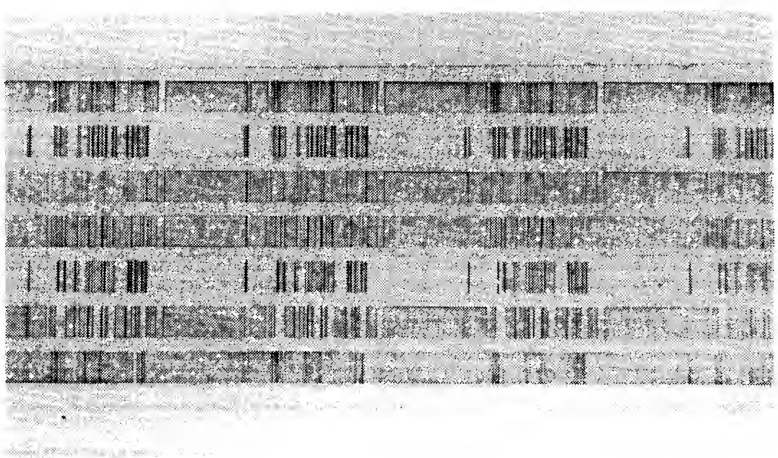

Figure 3-6. Data impression on transparent adhesive tape

Now the tape can be placed in a log book, or on a piece of white paper, along with other information about the data tape or tape deck.

The magnetic tape is not harmed by this procedure. But you must handle the tape carefully and thoroughly wipe it clean before winding it back on the reel. (See figure 3-7.)

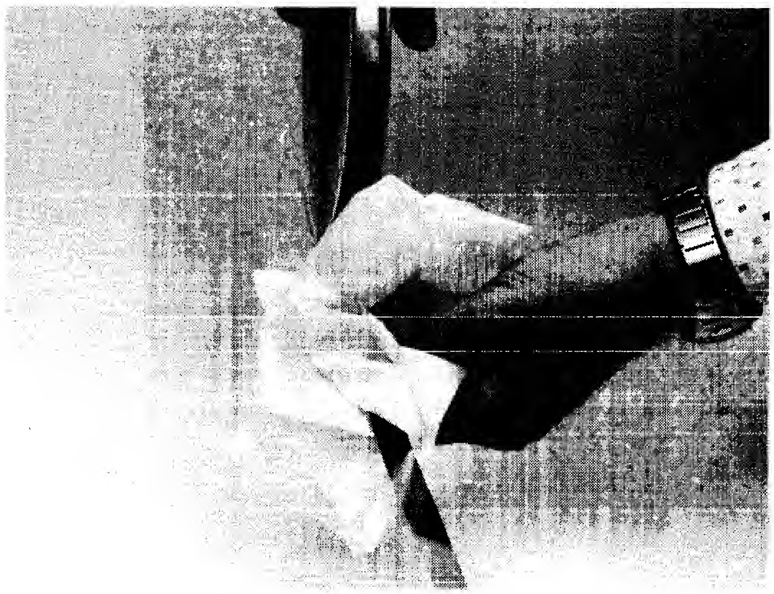

Figure 3-7. Wiping the magnetic tape clean

Error Detection Techniques

The reading and writing of data on magnetic tape is an electromechanical process. Unfortunately, activities that involve electronics and mechanics are never ideal and errors can occur. They may be caused by physical damage, dirt, changes in the supply voltage from the power company, mechanical wear, electrical noise, or any number of other sources. When you consider that a speck of dirt that is only .000140 inch can cause a bit of data to be dropped at 800 B.P.I., it is no wonder that several methods of detecting errors on tape have been developed.

Parity Checking

Most of these methods involve some form of parity checking, a process in which each group of bits is counted to see if there is an odd or even number of bits.

If odd parity is used, a parity bit will be added to the data only if the bit count is an even number. Adding the parity bit to the data will then make the count odd. In this manner, every time a group of data bits is read, there should be an odd number of bits. If the number is even, then you know that there is a data error. The parity bit does not function as part of the data, but is used only in the parity checking process. Normally, the parity bit is written at the same time that the data is written and is checked when the data is being read.

Even parity may be used, also. Even parity works the same way as odd parity, except that the parity bit is added whenever there is an odd number of data bits, so that the count of data bits plus parity bit will always be even.

Frame Parity

In actual practice, each frame of data on tape contains a parity bit. Therefore, a frame on a 7- or 9-track tape contains either 6 or 8 bits of data plus a parity bit. In phase encoding, group coded recording (GCR), and 7-track NRZI in binary mode, the frame, or vertical parity, is odd. In 7-track NRZI in BCD mode, the vertical parity is even.

Even with frame parity checking, there are many instances when a combination of bits in a frame can be in error, yet not cause a parity error. For example, suppose you are maintaining odd parity and have a frame with 5 bits. If 2 bits are dropped, you would still have an odd number of bits remaining and no errors would be detected. Because of this, frame parity checking is always used in conjunction with other methods of error detection.

Longitudinal Redundancy Check

One method used in conjunction with frame parity is the longitudinal redundancy check character or LRC. The LRC counts all of the bits in each track and, at the end of each record, adds a parity bit at the end of every track. The LRC is usually even parity. This way, if errors in a frame go undetected, they will probably cause an error in the LRC. (The LRC is found on NRZI tapes.)

Cyclic Redundancy Check

Another check character is the cyclic redundancy check character or CRC. The way in which this character is derived is quite different from how the simpler parity check characters come about. In obtaining the CRC, each frame of a record is put through a half adder register.

There are two advantages to using the CRC character. The technique used in its derivation nearly eliminates the possibility of undetected errors. The exact nature of simple data errors can be determined and, therefore, corrected. In practice, if a CRC error is detected, the data can be reread and, by using the vertical parity with the original CRC, simple bit errors can be corrected in the tape control unit. Nine-track NRZI and GCR tapes contain CRCs.

Error Correction Code

GCR tapes also contain auxiliary CRCs and error-correction code characters, or ECCs. These characters are derived in the manner that CRC characters are, but with different mathematical formulas. They are also used in error detection and correction.

PE Error Detection

In phase encoded tapes, an error detection and correction scheme is inherent in the phase encoding recording method.

As you may recall from the discussion of the phase encoding recording method, each bit of data is considered to be in the middle of an imaginary cell. Because of the recording technique, flux changes may be necessary between the cells. These flux changes between the cells can predict what will appear in the next cell. Following is a summary of this:

1. A positive pulse in the middle of a cell indicates a 1; a negative pulse indicates a 0.

2. A positive pulse at the junction of two cells indicates that both cells contain 0s; conversely, a negative pulse indicates that both cells contain 1s.

3. The absence of a pulse at the junction of two cells indicates that the cells contain opposite binary bit values.

The presence or absence of bits between cells can be used to predict the next bit. Thus, this between-cells activity, when combined with the frame parity, can help detect simple data errors and correct them while the tape is being read. This is sometimes known as "on-the-fly" error correction.

File Protection

File protection is a method by which reels of tape are protected from accidental overwriting.

Write Ring

Each reel has a circular groove cut into it, on the side that is mounted next to the tape drive. In one form of file protection, a circular plastic ring approximately .25 inches thick and 4.25 inches in diameter, is snapped into the groove to allow writing on the tape.

Each tape drive has a small pin that is spring-loaded to extend into the groove of a reel when it is mounted. This pin is attached to a microswitch inside the tape drive; the microswitch is connected to the read/write logic and will disallow writing on the tape if it does not detect the plastic ring.

This is an important feature, because an accidental erasure of a tape could result in the loss of many hours as the destroyed records are recreated. Such a loss could be brought about unintentionally by a software or hardware malfunction or by a case in which the computer is told to follow the wrong program. With the write ring removed, these types of accidental erasure cannot occur.

All newly purchased tapes come with a "write ring" installed. When it has been determined that no more records will be written on a tape, the write ring is removed. It can be inserted later to write new records on these used tapes.

The presence or absence of the write ring has no effect on the reading of the tape. This ring is sometimes called the "file protect ring."

Tape Labels

Another form of file protection uses software. Whenever the operating system creates a new tape, it writes at least two special records at the beginning, which are referred to as "tape labels." The first record is called the VOLSER or VOL1 (volume serial number). It is a unique serial number for that tape. The VOL1 record is usually 80 characters long, but only 6 characters are used for the serial number. The rest may be blank. The next record is called HDR1 (header 1) and contains 80 characters identifying such information as creation date, system codes, version number, etc.

The purpose of these records is to identify the files on a tape. When a job is running under an operating system, it may call for a specific tape to be mounted on a specific drive. The operating system will always check the first two records (VOL1 and HDR1) of a newly mounted tape and thus compare it with the request. If the operator accidentally mounts the wrong tape, the operating system will not use it, and will inform the operator of the error.

Tape Handling (Text)

As magnetic tape recording densities are increased to cope with modern data processing technologies, greater demands are placed on magnetic tape performance. Using the finest magnetic tape with the highest quality tape transport equipment is a requirement to provide maximum results in computer recording. However, other factors also affect recording quality. Improper handling and storage of tapes can result in a variety of operational problems, any of which can cause signal distortion or complete dropout of recorded data. Since information may be written only a few thousandths of an inch from the tape's edge, simple but effective precautions must be taken to avoid data loss through damage.

Aside from the specific tape loading procedures found in the instruction manual for the tape transport being used, there are certain tape handling precautions which should always be observed.

Irregular Wind

High-speed tape winding operations cause air to be trapped between the tape layers; this tends to cause the tape to stack irregularly on the reel. When this occurs, the tape edges protrude slightly, forming what appears to be an irregular surface when viewed through the reel flange cutouts. This condition does not hamper tape performance, but does require careful reel handling to avoid squeezing the flanges into contact with the tape edges. Handle tape reels at the central hub area whenever possible.

Reel Care and Handling

Extra care should be taken when removing a tape from the transport hub. Avoid the tendency to squeeze the reel flanges together when pulling a reel off the hub.

Improper seating of reels, improper threading of tape on the transport, and improper handling of reels can cause stretched, wrinkled, or creased edges. A wavy-edged condition prevents proper tape-to-head contact and results in serious loss of signal amplitude and intermittent errors. The wrinkled edges present a stretched appearance and normally do not lie in close contact with a flat surface. Reel warpage or improper insertion of the write ring can also cause this form of edge damage. Other kinds of tape edge damage include nicks and creases due to squeezed flanges.

Precautions should be taken when a reel of tape is removed from its canister. The end of the tape tends to unwind from the reel, thus exposing many feet of tape. Although the first 10 to 15 feet are not used for recording, they are threaded through the tape transport guides. If the end of the tape is allowed to touch the floor or come in contact with a dirty surface, dust and dirt can adhere to the tape, and can be transferred to the transport guides and onto the heads.

Some tape transports are self-loading. These autoloading drives will accept a tape with an outside protective strap that is opened by the drive automatically. This eliminates the probability of any tape unwinding from the reel while it is being handled.

Tape Storage

Tapes should be stored on their edge and should not be exposed to magnetic fields. If a tape is stored at an abnormal temperature or humidity, it should not be used until it has stabilized at room temperature.

Tape Marker Lab

Introduction

In this activity, you will place a tape marker on a reel of tape.

Materials Required

Reel of magnetic tape
Spool of magnetic tape markers (IBM #352407)

Lab Procedures

1. Unroll 10 to 15 feet of tape from the reel. Do not allow the tape to touch the floor or become contaminated. (See figure 3-8.)

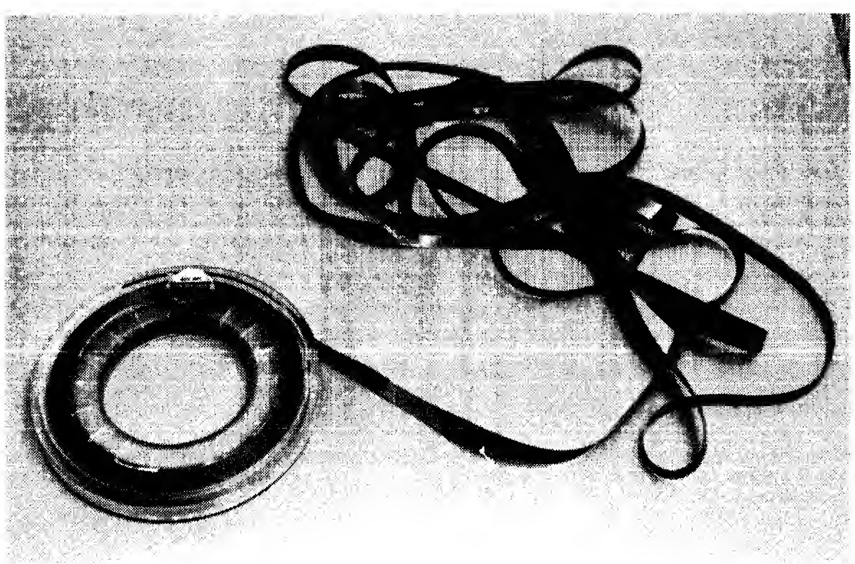

Figure 3-8. Tape unrolled on table

2. Hold the spool in your left hand and grasp the carrier tape on the back with your right hand. (See figure 3-9.)

Figure 3-9. Holding and grasping

3. Be sure the carrier tape is under the staple, then carefully pull the carrier tape straight down until one of the markers is approximately half removed. (See figure 3-10.)

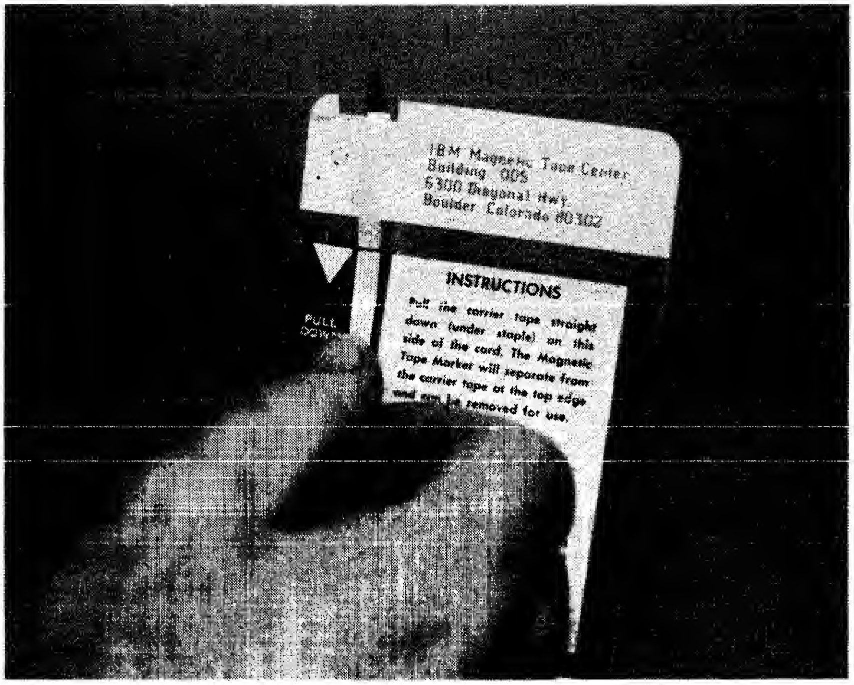

Figure 3-10. Back of card, hand pulling carrier tape

4. On the shiny side of the tape, carefully align the marker to be slightly in from the edge. (See figure 3-11.)

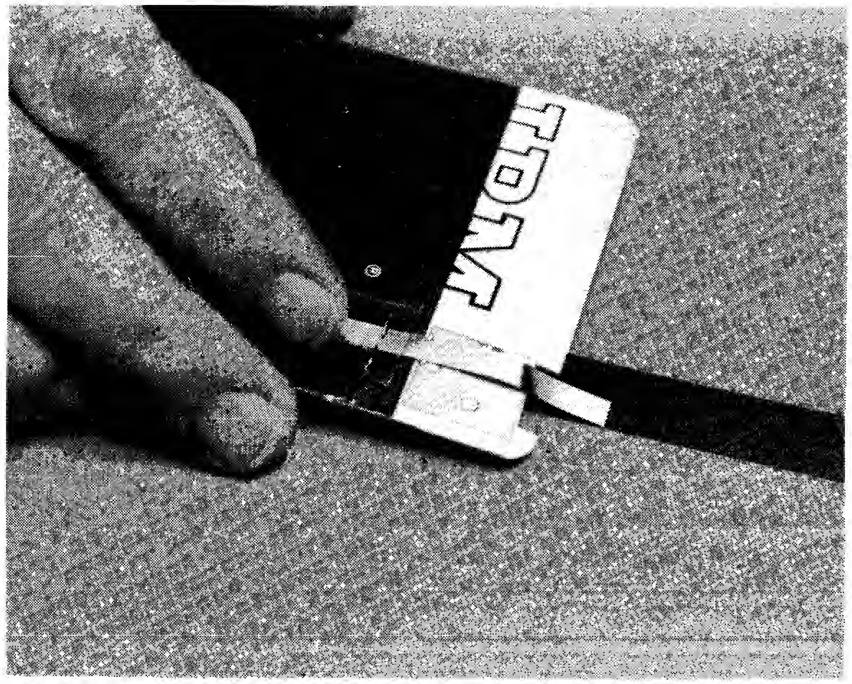

Figure 3-11. Placing marker over tape

5. Maintain alignment and press the marker into place. (See figure 3-12.) Carefully slide the applicator to the left to completely remove the marker from the carrier tape.

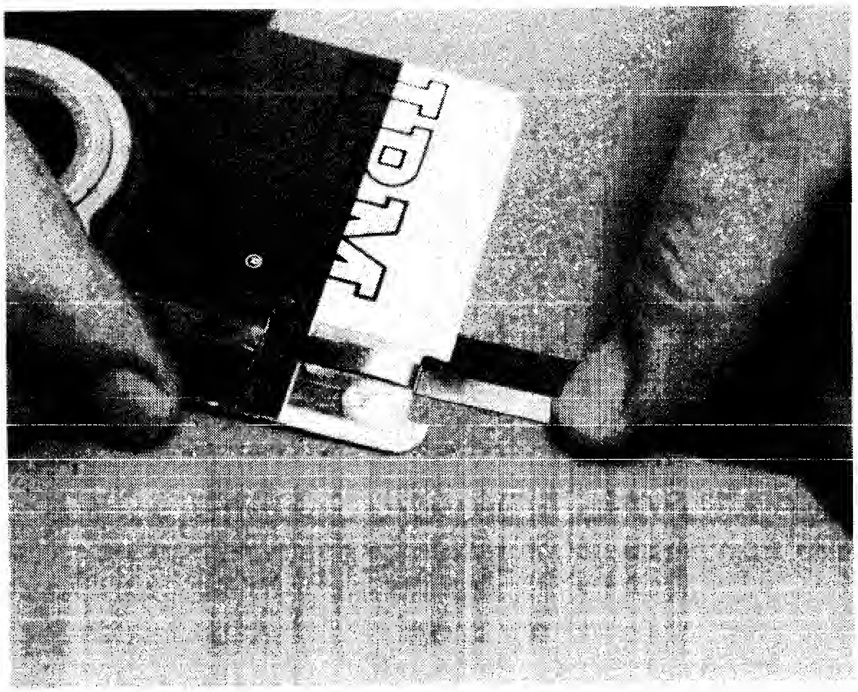

Figure 3-12. Pressing marker into place

6. Smooth out any bumps with your fingernail. (See figure 3-13.)

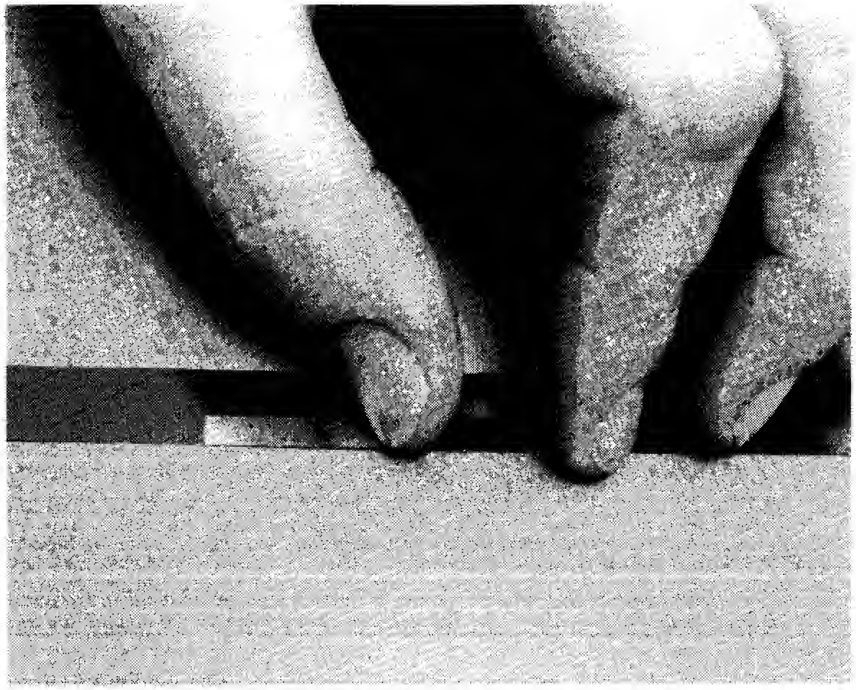

Figure 3-13. Rubbing with fingernail

7. Carefully inspect the marker to ensure that it does not go past the edge. (See figure 3-14.) If it does, it should be removed and disregarded and a new marker installed.

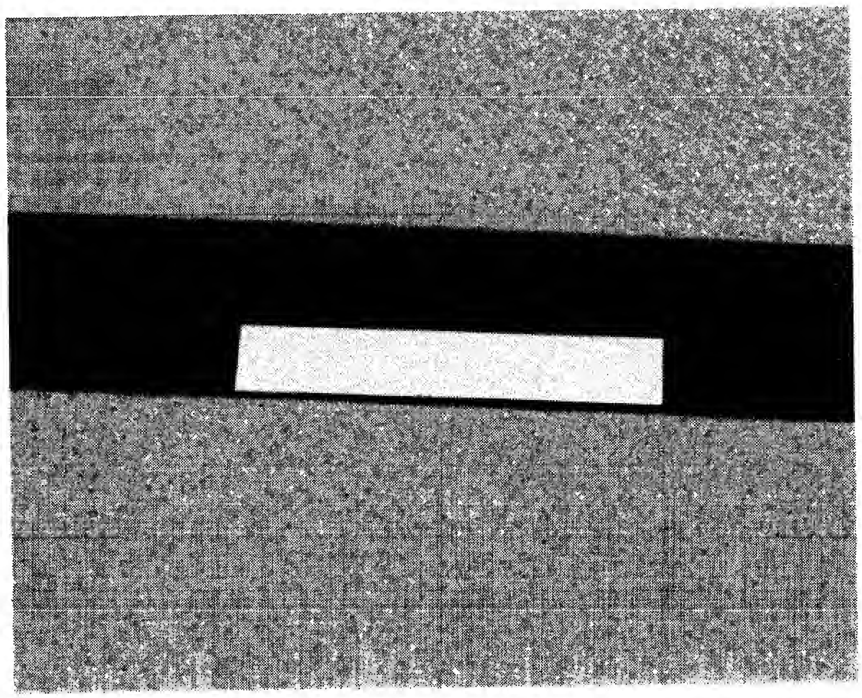

Figure 3-14. Marker properly installed

8. When you are satisfied that the marker has been properly installed, remove it and wind the tape back on the reel. (See figure 3-15.)

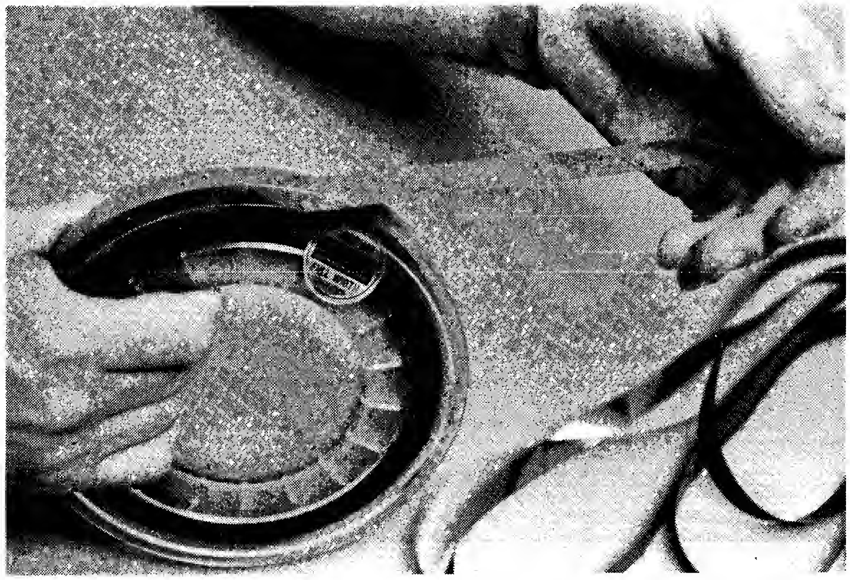

Figure 3-15. Rolling up tape on reel

Block 4

Introduction to Tape Transports

Transport Architecture

There are many manufacturers of tape transports, and each produces a series of models. Each model is unique in appearance and operation, but all have certain common characteristics.

Tape Handling

The tape handling methods employed by all tape transports are very similar. Each tape drive has two hubs on which the supply (full) reel and take-up (empty) reel are mounted. A capstan is also necessary to pull the tape across the heads.

Transports may have one or two capstans. The newer models tend to have the single capstan drive. The single capstan must handle both forward and reverse operations. This simplifies the mechanics, but complicates the capstan control circuits, because it requires that the capstan motor be reversible.

Older tape transports use pinch rollers with a continually rotating capstan to move tape, while newer models use a capstan with a vacuum applied to pull the tape onto it or air pressure to push the tape away. Thus are start and stop operations accomplished.

Tape-buffering is usually done with vacuum columns. This method of buffering has not changed much over the years. There are two methods of detecting tape in the columns; one method uses vacuum sensors and the other uses photosensors and a light source.

Reel and capstan control circuits have become more accurate as solid-state devices and have become more sophisticated and less expensive.

Many tape transports have to be threaded manually, but the newer, more expensive transports are threaded automatically.

Most transports have the supply (file) reel on the right and the take-up (machine) reel on the left; however, some transports, such as the IBM 2400 series drive, have reels arranged in the opposite way.

Tape Loading Methods

Tape loading starts with the mounting of a tape on the transport. This is accomplished by the following steps:

1. Select the proper tape and transport.

2. Open the window and mount the tape.

3. Manually thread the tape onto the take-up reel.

4. Wind several turns of tape onto the take-up reel to insure no slippage.

5. Close the window and press the load button.

Now the tape transport begins the load sequence. The tape will be moved forward slightly and then be pulled into the vacuum columns. The tape will be moved forward again until the BOT is sensed.

If the vacuum columns have tape in them, and the BOT is sensed to be in the correct position, the tape is at the load point. If no errors are indicated at this time, the READY switch may be pressed to indicate that the tape unit is loaded and ready. The ready status is sent to the controller and/or the computer.

Some of the newer drives or units have autoloading, which eliminates the manual threading of tape by the operator.

Autoload Devices

Autoload, a feature incorporated on some tape drives, automatically threads and loads tape. The actual autoloading is primarily performed by a combination of vacuum and air pressure being diverted to the right place at the right time. Autoloading can be performed on standard reels of tape or on tapes enclosed in cartridges. Cartridge tapes mounted on the tape drive are opened by the drive during an autoload sequence.

Autoloading consists of the following steps:

1. The operator places the tape cartridge on the file hub and pushes the LOAD switch.
2. After a short delay (two to three seconds) to allow the buildup of vacuum/pressure, the hub secures the reel to itself and opens the cartridge.
3. The window is now closed.
4. The vacuum is applied to the take-up reel hub.
5. The file (supply) reel is slowly rotated clockwise.
6. Air jets blow the tape through the tape guides and onto the take-up reel.
7. The tape is wrapped around the slowly rotating take-up reel because of the vacuum in the hub. The rotation is clockwise.
8. After two to three seconds, a vacuum sensor determines if tape is on the take-up reel.
9. If no tape is sensed, the load sequence is aborted and an unload sequence is started.
10. If a tape cartridge is mounted, a second attempt is made to load the tape. A third attempt is not made.
11. If the tape sensors indicate that threading was successful, the tape continues to be wound on the take-up reel until the BOT marker is sensed.
12. When the BOT marker is sensed, air pressure is removed from the tape guides and a vacuum is applied to the vacuum columns.
13. The left (take-up) reel is now reversed and will rotate counterclockwise while the right (file) reel continues to rotate clockwise.

14. Tape is pulled into the vacuum columns by the vacuum present, until the sensors indicate the presence of tape.

15. The reels are now under the control of the sensors in the vacuum columns.

16. The tape drive now moves the tape forward for about one second.

17. Then the tape drive moves the tape in the reverse direction until the BOT is detected.

18. A successful load is indicated by a load point light.

Most autoloading tape drives load tape in this manner, while the processes of others may vary slightly.

Manual Loading (Text)

It is the responsibility of operators and customer engineers to mount and thread the tape on tape transports. Care must be taken to prevent damage to the tape during threading.

Before tape loading can commence, the tape transport's window must be opened. Some drives have power windows, while others have ones which must be opened and closed manually. The power-operated windows usually have a bar across the top. Pushing up on this bar opens the window. The bar is specially positioned for safety reasons. If your hand is in the window opening and the window closes, your hand will be pushed up against the bar, tripping a switch which will open the window.

Once the window is opened, the reel is placed on the hub of the transport by pushing only on the reel hub. Do not push on the reel flanges, since doing so may cause damage to the tape edges.

Be sure the reel's write ring slot is next to the transport.

Now the reel must be secured to the hub; there are three common methods of doing so. One method uses a large knob on the hub that is tightened by turning it clockwise. This causes the hub to expand and grip the reel. Another method uses a flange extending from the center of the hub which is pushed in to expand the hub and secure the reel. The third method has no manual control on the hub; the hub expands by vacuum/pressure to grip the reel when the load button is pushed. This last method is used only on autoloading tape drives.

It is important that the reel is flush with the drive before the reel is latched to the hub; otherwise, the reel will wobble and damage the tape, and possibly the drive, if the reel falls off.

After the reel is secured to the hub, four to five feet of tape must be unwound in order to thread the tape through the tape guides and onto the take-up reel. The proper procedure for unwinding tape from the supply reel is to rotate the reel hub by hand.

Some tape transports are constructed so that their reel motor brakes are on even when the drive is not in use. To remove tape, you must release these brakes. If you work with a transport like this, you will find a pushbutton near the reels which, when pressed, momentarily releases the brakes while turning the reel.

After four to five feet of tape have been unwound, the tape must be threaded through the tape guides and onto the take-up reel. This threading procedure varies slighty with different models.

Threading may be from left to right or right to left, depending on which side the supply reel is on. (See figure 4-1.)

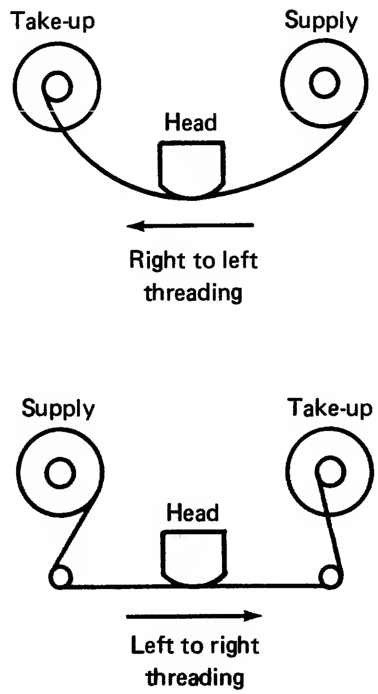

Figure 4-1. Tape threading

Tape must be wound several turns onto the take-up reel to insure that no slippage occurs when the load operation takes place. The reel brakes may have to be released to do this.

After the tape is threaded, the window must be closed. If the drive has a power window, pushing the load button will close it. After the load button is pushed, the transport will begin the load sequence.

A vacuum will be applied to the vacuum buffer columns while the supply reel is rotated clockwise and the take-up reel is rotated counterclockwise. The tape is pulled into the vacuum columns, where it is sensed by the vacuum column sensors. The transport now moves the tape in reverse until the BOT is detected.

After BOT is detected, and the loop sensors have determined that the tape is correctly positioned within the loop, the load point indicator will light.

At this time, the READY switch may be pressed which will indicate to the system that the unit is loaded and ready. This is also known as "on-line" mode.

On most units, the READY switch may be pressed any time after the LOAD switch has been energized; however, a READY condition and/or status will occur only after the tape loops have been loaded and BOT (load point) has been reached.

Block 5

Tape Transport Mechanics

Tape Motion

In less expensive audiotape decks, the mechanical drive for the supply reel, take-up reel, and capstan is provided by a single motor with the aid of several belts and pulleys. More expensive audiotape decks employ three separate motors for more accurate control of the capstan and reels.

Digital tape transports must have very accurate control of the capstan and reels.

Audiotape decks record and play back at fairly slow speeds compared to those of digital tape drives; therefore, you can stop audiotape decks quickly simply by disengaging the capstan and applying mechanical or electrical braking to the reels.

This method cannot be used on digital transports because of larger reels, higher operating speeds, and start/stop times that must occur in milliseconds.

The solution to the problem of reducing start/stop time is to separate the starting and stopping of the tape from the starting and stopping of the reel motors. In other words, the movement of tape across the head should be independent of reel control.

Figure 5-1 is a typical tape transport, showing the major components of a tape drive.

The supply and take-up reel motors are controlled by a closed loop servo. A closed loop servo consists of one or more sensors which control a mechanical action. For example, a house may have a heating system which is a closed loop servo. When there is not enough heat, the thermostat (sensor) contacts close and the heating system comes on. When enough heat reaches the thermostat, the contacts open; this turns off the heat.

Looking at figure 5-1, you will find that the "thermostat" in this case is actually the loop sensing photocells, and the "furnace" is the reel motor.

When tape is pulled out of the right-hand vacuum storage column by the capstan, the upper column sensor photocell is illuminated and signals the supply reel motor to turn clockwise. This action puts more tape in the column until the sensor is covered. When the photocell no longer has light on it, the reel motor is signaled to stop.

If tape is moved in reverse, the bottom sensor is covered when the tape gets longer. This signals the supply reel motor to turn counterclockwise, which, in turn, removes tape from the column until the sensor is uncovered.

The take-up reel motor is controlled in the same manner as the other column and its sensors.

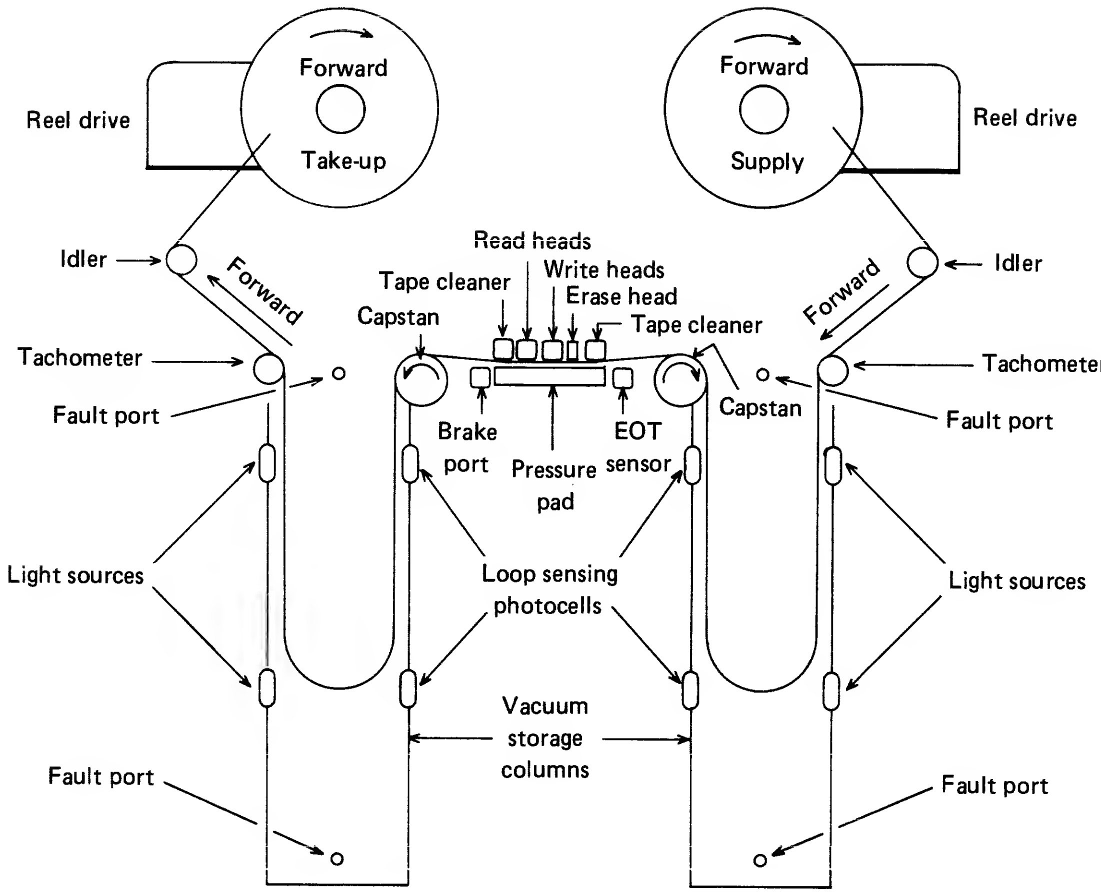

Figure 5-1. Tape transport

Tape Brakes

Tape movement across a head must be precise: 1) when tape movement is initiated, the acceleration must be rapid; 2) after the tape speed is attained, it must remain constant; and, 3) the stop time must be brief.

Start and stop time must occur within the record gap, which, if you recall, is about 0.6 inch. If the drive is reading or writing a record, it must stop at the end of the record, preferably in or near the center of the gap. This means it must stop 0.3 inch beyond the end of the record.

If tape drives move tape at speeds up to 200 inches per second, then simple mathematics gives the stop time as follows: 200 inches per second can be converted to 0.2 inch per millisecond. Therefore, to stop within 0.3 inch, the tape must be stopped in 1.5 milliseconds. One hundred inches per second would allow twice the stopping time, or 3.0 milliseconds.

Simply to disengage the capstan would result in continued tape motion, due to the inertia of the tape. To prevent this additional motion, brakes are used to stop the tape.

Types of Brakes

There are three common types of brakes used: mechanical brake pad, pinch roller brake, and vacuum brake.

Mechanical Brake Pad

The mechanical brake pad (see figure 5-2) consists of a solenoid, armature, and pad. When the solenoid is energized, the armature is pulled toward the solenoid. This squeezes the tape between the hard rubber brake and the brake pad.

Two brake assemblies (see figure 5-3) are used if the tape transport has forward and reverse operations. The mechanical brake pad method is only effective at relatively slow tape speeds.

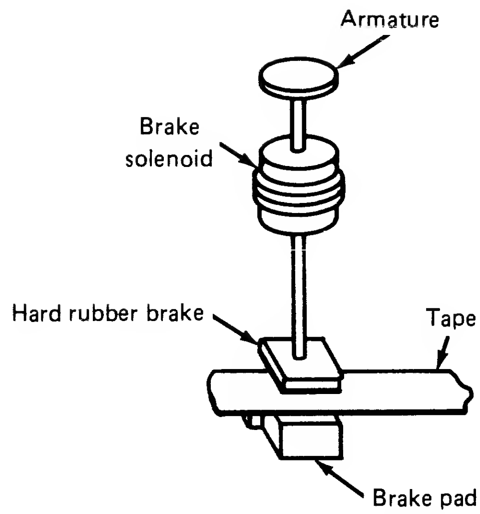

Figure 5-2. Mechanical tape brake mechanism

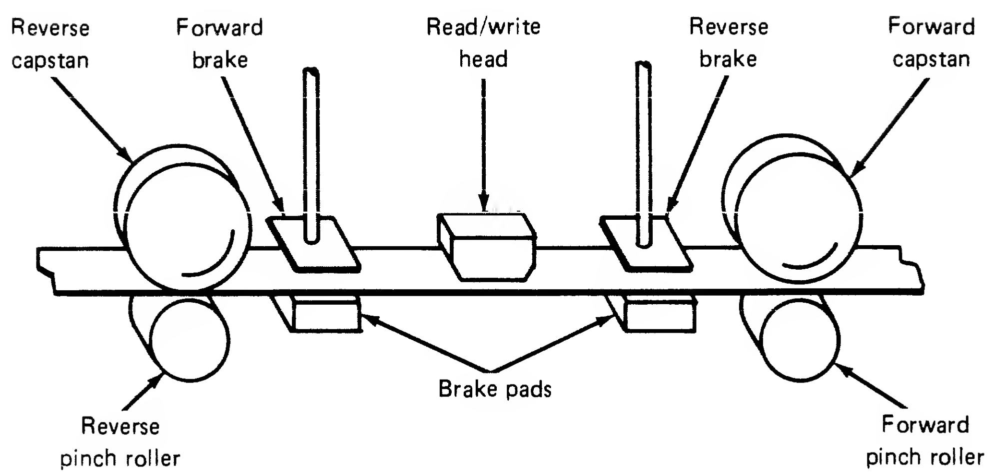

Figure 5-3. Arrangement of tape drive and tape brake mechanisms

Pinch Roller Brake

A pinch roller brake is usually used with pinch roller capstans. (See figure 5-4.)

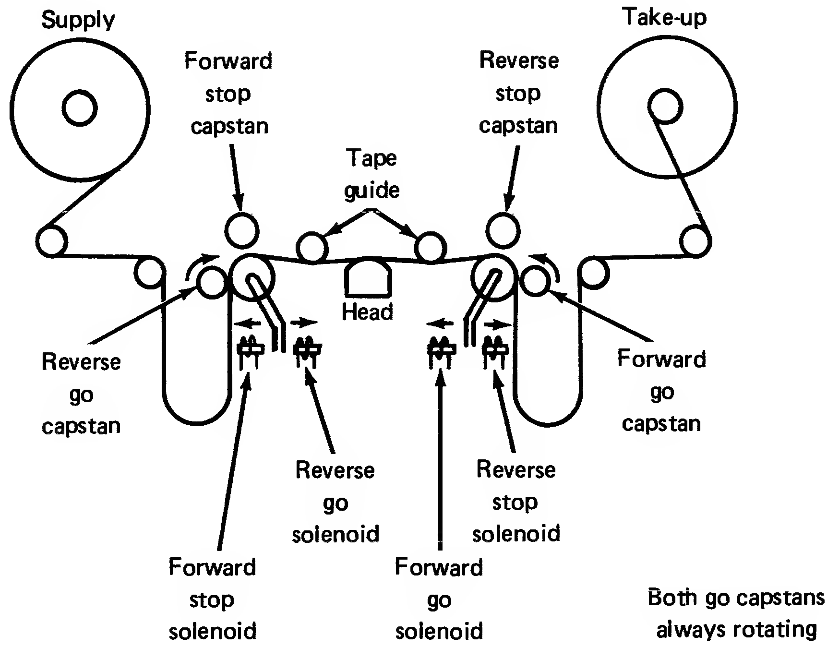

Figure 5-4. Pinch roller brake

Certain types of tape servo systems use a "prolay," or device that starts and stops tape motion. The prolay (see figure 5-5) has three positions: stop, neutral, and go. When the forward go solenoid is energized (see figure 5-4), the prolay is pulled against the capstan, which is constantly rotating forward. When the tape transport is signaled to stop, the forward go solenoid is deenergized and the prolay returns to neutral. At the same time, the prolay in figure 5-5, which was in neutral, has its forward stop solenoid energized. This forces the tape against the fixed forward stop capstan, thus stopping the tape.

Notice that, during forward start/stop operations, the left-hand prolay moves between the neutral and stop positions, while the right-hand prolay moves between the neutral and go positions. This short prolay travel results in rapid start and stop times.

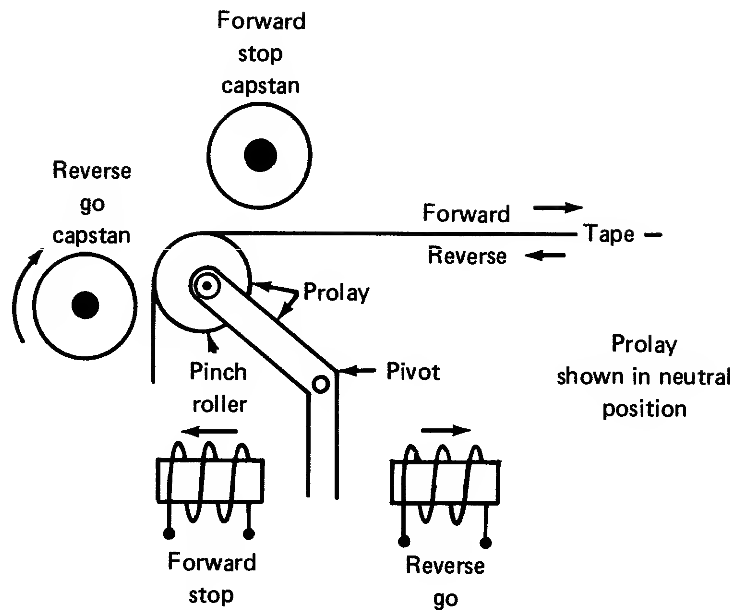

Figure 5-5. Prolay operation

If the tape transport were signaled to read backwards, the left-hand prolay would move between neutral and go, while the right-hand prolay would move between neutral and stop.

Vacuum Brake

The third type of brake used is the vacuum brake. (See figure 5-6.) It consists of a hollow chamber with many holes in the surface. The tape passes over this surface during read and write operations. This hollow chamber has two hoses connected to it, and a shaft extending out of the side. One of the hoses has a vacuum applied to it, and the other has air pressure. The shaft is connected to a selector valve inside the chamber to select either vacuum or air pressure.

As the tape passes over the surface of the hollow chamber, it is pushed away by air pressure. When the tape transport is signaled to stop, a solenoid is energized, which moves the shaft to select vacuum instead of air pressure. The vacuum pulls the tape firmly against the vacuum brake and stops the tape.

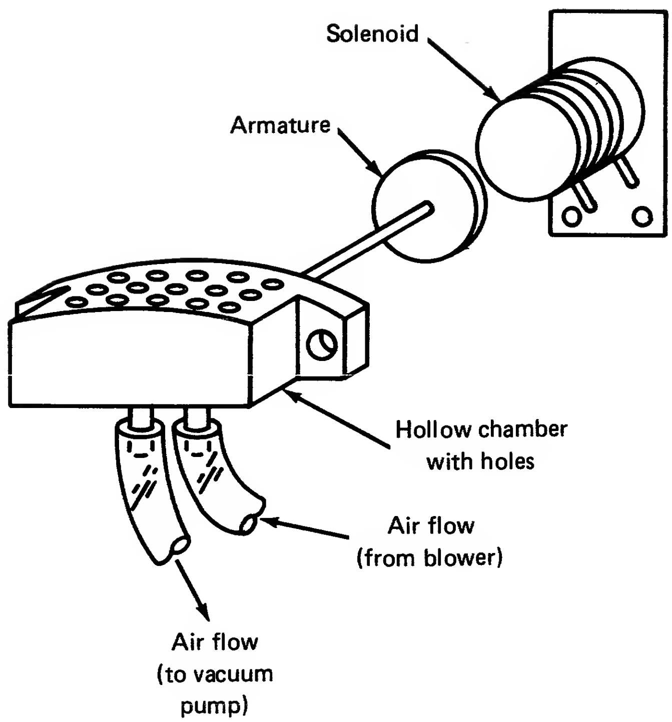

Figure 5-6. Vacuum brake

This form of braking is simpler, faster, and easier to maintain than the other two. Vacuum brakes are commonly used today in transports requiring some form of brakes.

Pinch Rollers

This activity describes the pinch roller capstan drive and gives additional information on the mechanics of prolay positioning.

Figure 5-7 shows the three possible positions of the prolay: neutral, reverse go, and forward stop.

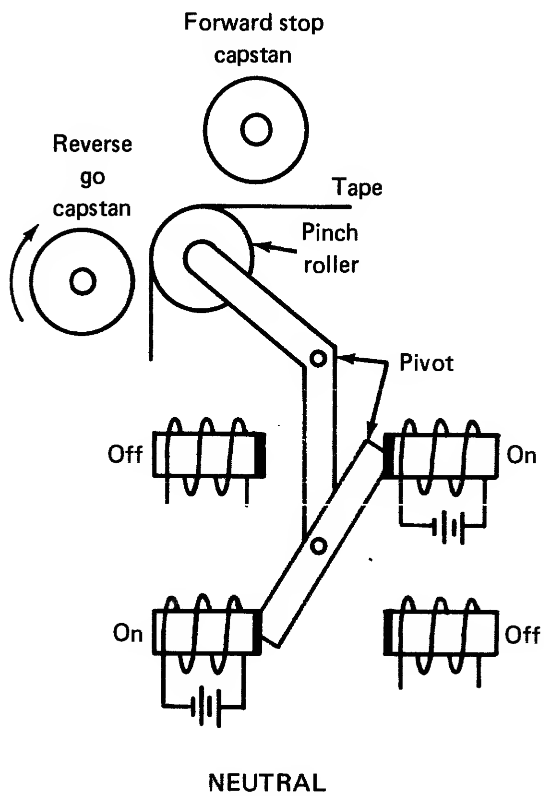

Figure 5-7. Positions of prolay

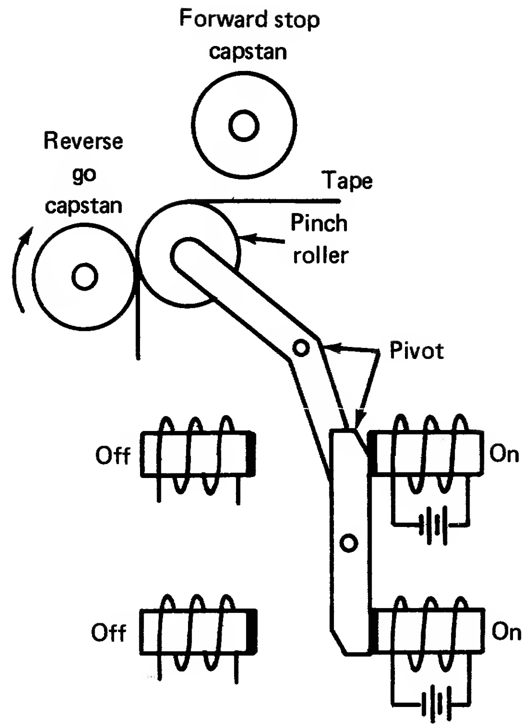

REVERSE GO

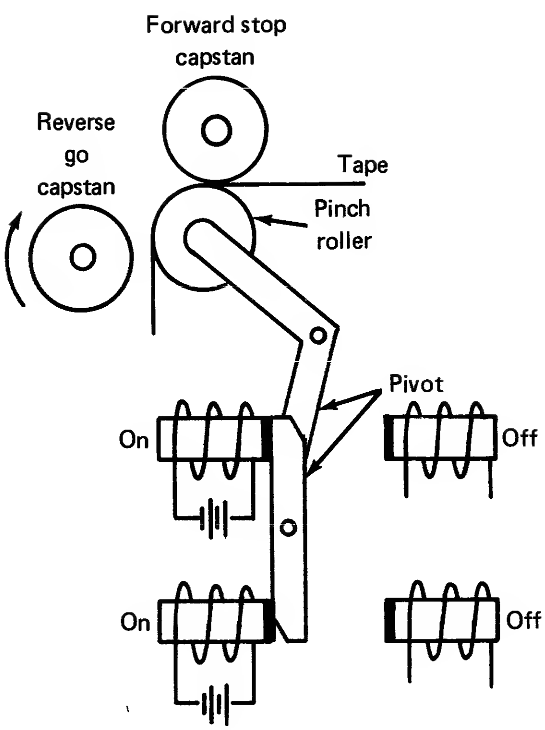

FORWARD STOP

Figure 5-7 – Continued

The travel distances in the drawings are exaggerated. The actual travel distance between the solenoids is only about .03 inch; however, the arm extending to the pinch roller is about two to three times longer than the arm extending down between the solenoids. This causes the pinch roller to travel about .1 inch between the stop and go capstans. The pinch roller is often called the "prolay idler."

Figure 5-7 shows only the left-hand prolay. The right-hand prolay is a mirror image, but is used for reverse stop and forward go.

Both forward go and reverse go capstans rotate continuously after the tape is loaded. Both capstans can be driven with one motor, as shown in figure 5-8.

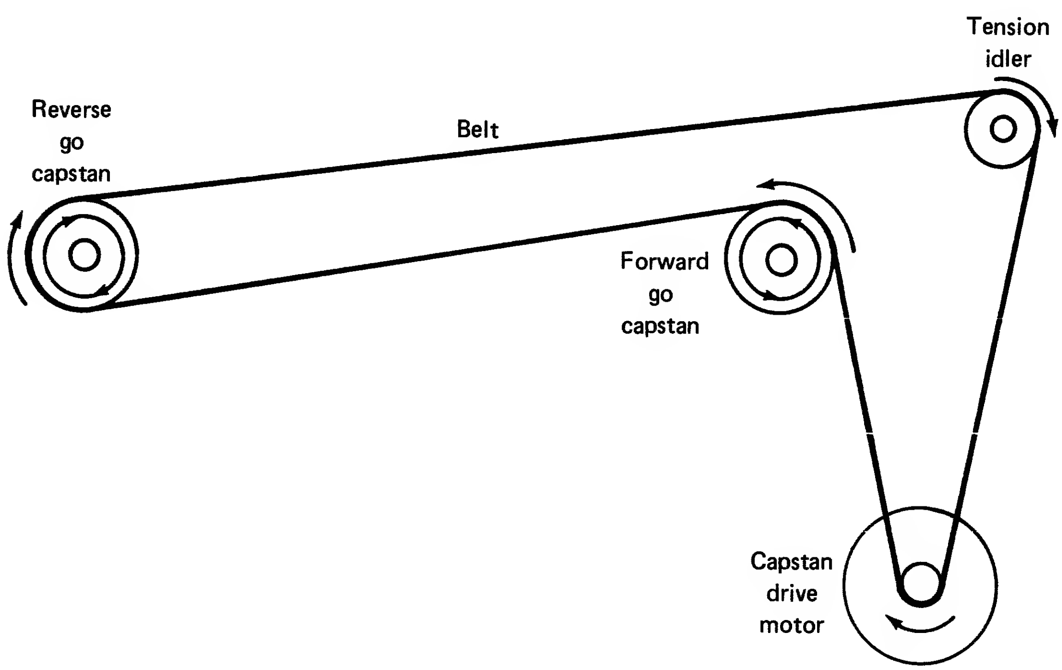

Figure 5-8. Single motor capstan drive

Better drives use a separate motor for each capstan, thus eliminating the belt and idler. This increases the reliability of the tape transport, because belts are always a source of trouble.

The tape path, after loading, is shown in figure 5-9.

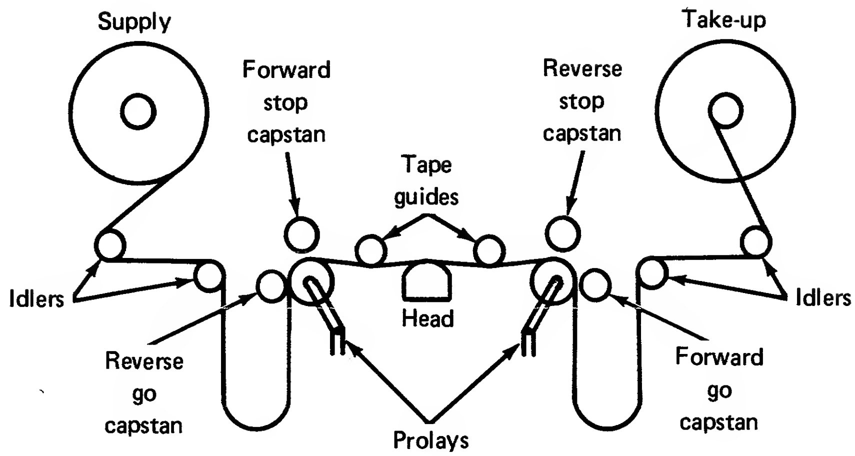

Figure 5-9. Tape path

Notice that the go capstans are actually in the loop vacuum columns. If the tape is unloaded, these capstans will be in the way. This problem is solved by adding a solenoid that retracts the capstans so that they are flush with the tape transport whenever the tape is being loaded or unloaded.

When the capstans are retracted, they trip a switch which removes power from the capstans and allows a load or unload sequence to start.

The tape guides (see figure 5-9) also move up and down. They are moved up whenever the tape is unloaded, making it easier to thread the tape. This type of transport must be threaded manually.

Most of the problems encountered with this method of tape handling are due to lack of lubrication and buildup of dirt. The stop and go capstans must be kept clean to prevent the tape from slipping.

Vacuum Capstans

A vacuum capstan consists of a hollow wheel about three to four inches in diameter and about .75 inch thick. (See figure 5-10.) Its surface has many holes or slots through which air pressure or a vacuum may be applied.

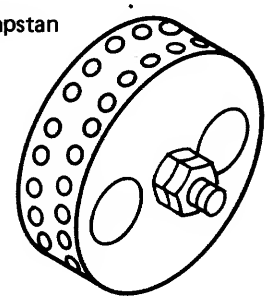

Figure 5-10. Vacuum capstan

Tape transports may employ vacuum capstans in one of two ways. They are called the "single capstan" and "dual capstan" drives. First, consider the dual capstan drive, which is the older version.

Figure 5-11 is a simplified drawing showing the positions of the capstans. Note that dual vacuum capstan drives also use vacuum brakes; single capstan drives do not.

When a forward read or write operation is signaled, air pressure is applied to the brakes, thus pushing the tape away from the brake surface. At the same time, a vacuum is applied to the forward capstan. Both capstans will be rotating whenever the tape is loaded.

With a vacuum applied to the forward capstan only, the tape is pulled onto the surface of the rotating capstan and is accelerated to read/write speed.

When the drive is signaled to stop, the vacuum is removed from the capstan and air pressure is applied. At the same time, air pressure is removed from the brakes and a vacuum is applied. The reverse capstan will always have air pressure applied unless a reverse read (read backward) operation is started.

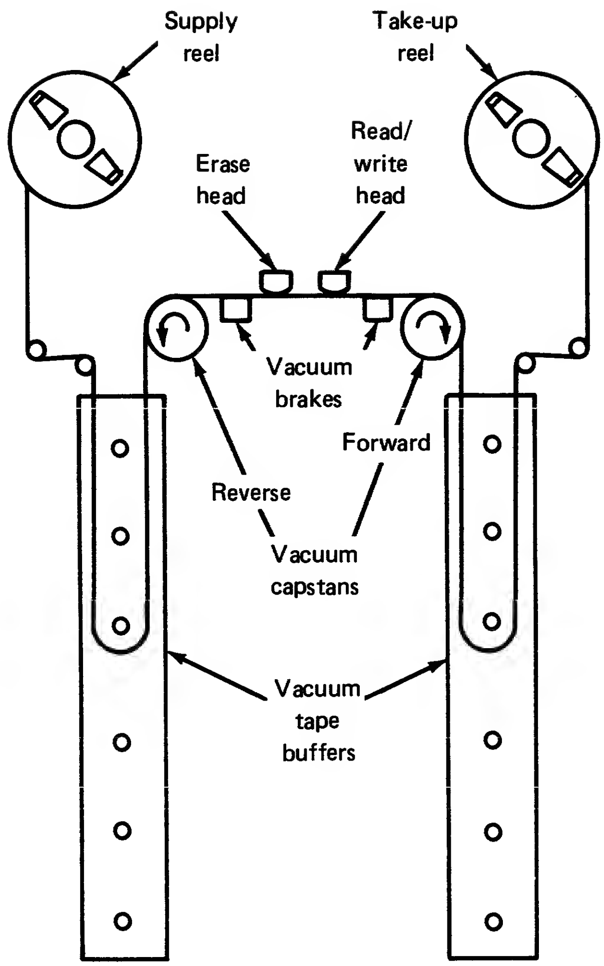

Figure 5-11. Capstan positioning

The switching between air pressure and vacuum is controlled by a valve. These valves may use a high-speed voice coil movement to decrease the switching time between vacuum and air pressure.

Vacuum is applied to both vacuum brakes for forward and reverse stop operations.

The advantage of the dual vacuum capstan drive over the pinch roller drive is that its use involves fewer moving parts, and thus it usually requires less maintenance.

The disadvantage is the complexity of the vacuum/air pressure control valves. These valves are subject to constant switching when the drive is started and stopped. They eventually become a source of trouble, failing to switch or reacting slowly to signals. Poor reaction time will cause the tape to overshoot the record gap when stopping or fail to be up to speed at the beginning of the record when starting.

The constant switching of the vacuum/air pressure valves was eliminated by the development of the single capstan drive, which has only one capstan with vacuum applied continuously after the tape is loaded.

The capstan rotates only when tape movement is required. It rotates counterclockwise for forward operations and clockwise for reverse operations.

To prevent tape slippage on the capstan during sudden start and stop operations, the tape is allowed contact with about three-fourths of the capstan surface. (See figure 5-12.)

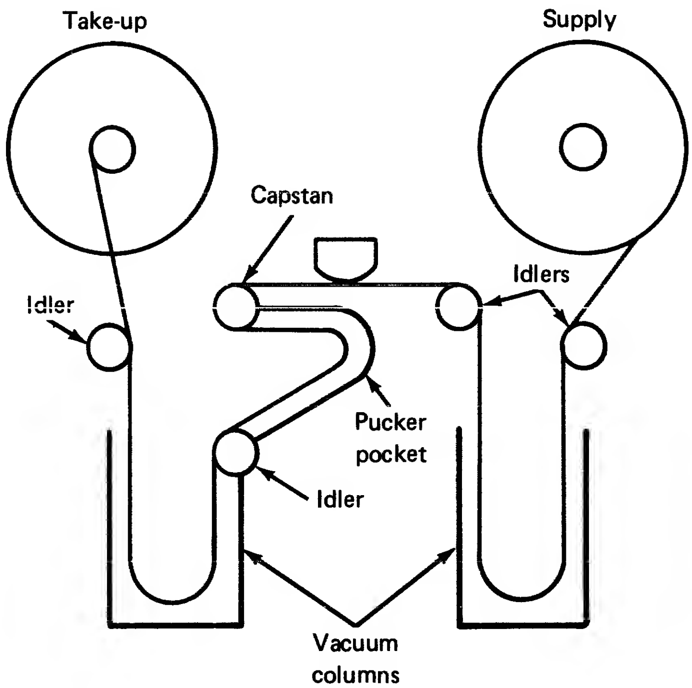

Figure 5-12. Capstan contact

This additional tape-to-capstan contact is accomplished with a pucker pocket, as shown. The pucker pocket has a vacuum applied constantly, like the vacuum columns.

The idlers are called "air bearings"; they are built like the capstan, but are smaller. Air pressure is applied to them to reduce the tape-to-idler friction.

The capstan motor is called a "printed circuit motor." It is controlled by a closed loop servo. The sensor is a digital tachometer, which senses the speed and direction of the motor and provides control signals to the motor drive circuits. This servo circuit can start and stop the capstan motor within a few milliseconds.

This closed loop servo system accelerates the motor quickly and maintains accurate control of its speed. It also eliminates the need for the constant switching between vacuum and air pressure that is required of continuously rotating dual capstan drives.

Tape Buffering

This learning activity will describe two methods of tape buffering. The need for buffering was discussed in the activity entitled Tape Motion.

There are two common methods for buffering tape: tension arm buffering and vacuum column buffering. Both use a closed loop servo.

Tension Arm Buffering

Tension arms are used on tape drives that operate at relatively slow speeds. They comprise a less expensive means to buffer tape, and are used more often on industrial tape drives than on computer tape drives.

Figure 5-13 shows a typical tension arm arrangement.

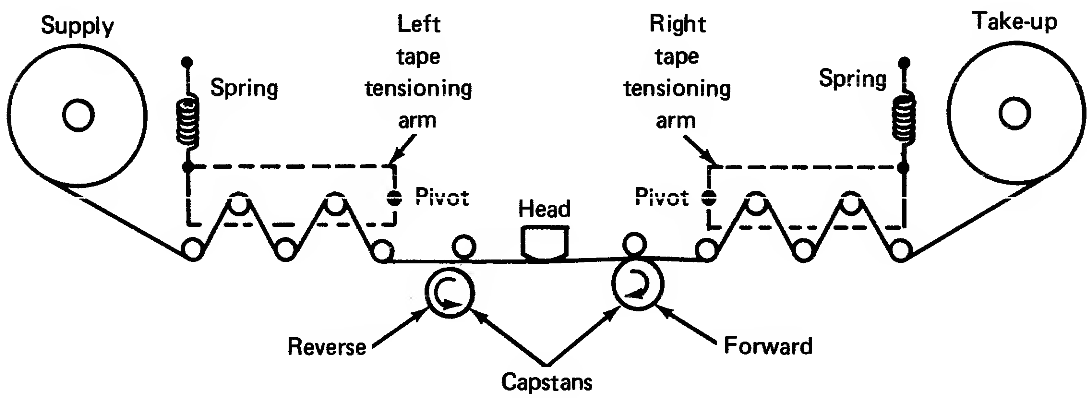

Figure 5-13. Typical tension arm arrangement

When forward tape motion occurs, the left-hand loops become smaller; this causes the left-hand tension arm to move down. This downward movement is sensed, amplified, and sent to the reel control circuits. This in turn causes the supply reel motor to turn counterclockwise until the loops are back to their normal length.

The forward motion of the tape also causes the right-hand loops to become longer, which causes the right-hand tape tension arm to move up. This movement is sensed, amplified, and sent to the take-up reel control circuits. The reel motor is then signaled to turn counterclockwise, which removes the additional tape in the loops. This pulls the tension arm down.

If either arm reaches its upper limit, the tape transport's reel motors and capstan motor are turned off. This happens if the tape breaks or the servos are out of adjustment.

Vacuum Column Buffering

Vacuum column buffering is the most common method used in computer tape transports. There are two types of vacuum column control: proportional and digital.

Proportional Control

Proportional control does not require reel brakes, but has a more complex servo. (See figure 5-14.)

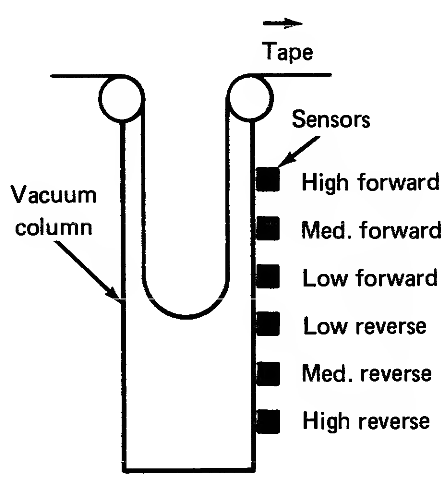

Figure 5-14. Vacuum column control (proportional)

When all the forward sensors detect tape and all the reverse sensors do not, the reel motor does not turn. As tape is pulled out of the vacuum column, the low forward sensor signals the reel motor to turn slowly and put more tape in the column. If tape is taken out of the column faster than the reel motor supplies tape, the medium sensor will signal the reel motor to turn faster.

This system works well on tape drives that can operate at two or more speeds or on drives that perform their high-speed rewind with the tape in the vacuum columns.

Digital Control

The digital control method uses only two sensors per column for reel control. (See figure 5-15.)

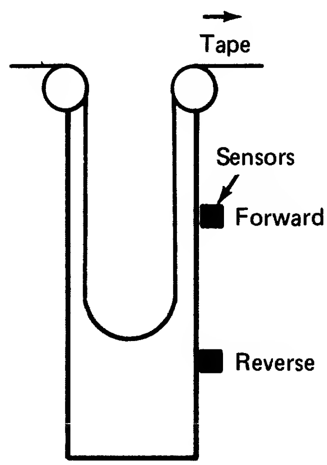

Figure 5-15. Vacuum column control (digital)

If the forward sensor does not detect tape, the reel motor is signaled to run full speed forward and dump tape into the column. If the reverse sensor detects tape, the reel motor is signaled to run full speed reverse and pull tape from the column.

When power is removed from a reel motor running full speed, it will coast and possibly dump in or remove more tape than it should. This will cause the reel motor to fluctuate between forward and reverse.

To solve this problem, reel brakes are used. If the forward sensor detects tape and the reverse sensor does not, the reel brake is applied.

Although this method simplifies the servo system, it requires additional circuits to turn on and off the brakes.

As soon as the reel motor power is removed, the brakes are applied. The type of brakes used by various drives will be discussed in another learning activity.

Pneumatic Systems

Almost all digital tape transports require a source of vacuum and air pressure. Even transports that utilize pinch roller servo systems need vacuum for their vacuum storage columns. Of course, transports that use vacuum capstans also need vacuum and pressure for the capstans and air bearings. It is the task of the pneumatic system to supply vacuum and pressure when needed. This activity introduces the areas in which pneumatic systems are required.

Vacuum Storage Columns

(Vacuum storage columns are referred to by several other names: tape buffer columns, vacuum buffer columns, vacuum columns, buffer columns, and loop box. All of these terms refer to the same thing.) Vacuum storage columns use low vacuum readings since only enough vacuum is required to hold the tape in the column and give it its characteristic loop shape.

Vacuum storage columns require vacuum pumps that can handle many variations in volume, because there are usually rapid changes in the positions of the tape loops and because the columns require a certain or constant volume. Therefore, a vacuum pump to serve vacuum columns must have high volume and a low vacuum reading.

Vacuum Capstan

Vacuum capstan servo systems require higher vacuum readings. This high vacuum is necessary to hold the tape firmly against the spinning capstan and against the tape brake. Because of the difference between the vacuum needed for vacuum storage columns and for servo systems, tape transports may use separate pumps for these two functions.

Tape Cleaners

Another use for vacuum is in the tape cleaners. These devices are hollow chambers with a slotted or perforated cover over which the oxide side of the tape rides. A small amount of vacuum is supplied to the cleaner so that loose particles of dirt and oxide are pulled from the tape before it passes under the head.

Plenums

Since servo systems switch from vacuum to pressure very rapidly and many times in succession, a pump connected directly to the servo system might not be able to keep up with the demand for vacuum and/or pressure. To avoid this problem, some pneumatic systems use plenums. A plenum is a chamber connected between the pump and the device using the vacuum or pressure. It simply holds a reservoir of vacuum or pressure from which the servo system can draw in periods of high demand.

Air Bearings

The same pump that creates vacuum also creates pressure. Pressure may be used in automatic tape loading systems, and is required for transports that use air bearings. Air bearings serve a function similar to idler rollers, but are slotted like vacuum capstans. When air pressure is applied to the slots, the tape is held off of the bearing in a manner similar to vacuum capstans. This provides an almost frictionless surface for tape to move over.

Pneumatic System Maintenance

Pneumatic systems frequently require significant preventive maintenance. The levels of vacuum and pressure must be periodically adjusted and filters that are on all air intakes must be cleaned regularly. In vacuum capstan servo systems, malfunctions in the pneumatic system may show up first as read/write errors on tape since changes in vacuum and pressure can have a significant influence on the time it takes to start and stop the tape between records.

Block 6

Tape Transport Mechanics Control

Reel Drive Control

To avoid stretching or breaking the tape by sudden reel motion or reel weight occurring in the tape transport, tape is fed into vacuum storage columns. The amount of tape in these columns is monitored, and when too little or too much tape is contained in a column, a signal is sent to the reel drive motor to feed out or take up tape. This activity describes how the motors are controlled in a typical tape transport.

Monitoring Devices

The two most common methods of monitoring the tape are vacuum switch and photosensor.

Vacuum Switch

The vacuum switch, as its name implies, operates with a vacuum supply. (See figure 6-1.)

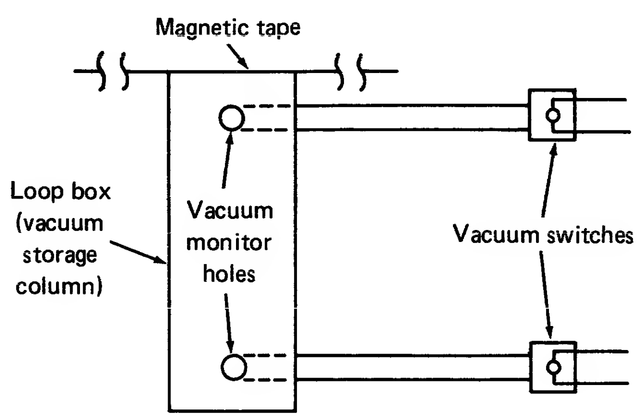

Figure 6-1. Vacuum switch configuration

When tape is manually loaded onto the tape transport, it covers the top of the loop box, power is applied, and the vacuum system turns on, thus creating a vacuum in the loop box. As the vacuum increases, the vacuum switches are activated via the vacuum in the vacuum monitor holes. The atmospheric pressure on top of the tape pushes the tape down and the vacuum in the column pulls on the tape; this action loads the tape in the loop box. When the tape passes across the top monitor hole, the pressure becomes atmospheric, causing the vacuum switch to become disabled.

Photosensor

The photosensor operation is similar, except that a light source and sensor replace the monitor hole. (See figure 6-2.)

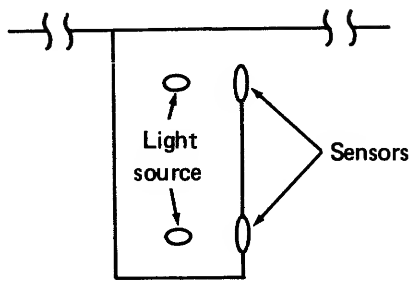

Figure 6-2. Photosensor configuration

As the tape passes in front of the sensors, light is interrupted and the sensors become disabled.

Operation of Monitoring Devices

The two reel motors (supply and take-up) are controlled by three signals from their respective vacuum storage column:

- Feed tape – this signal turns the supply reel clockwise (CW) and the take-up reel counterclockwise (CCW).
- Take-up tape – this signal turns the supply reel counterclockwise (CCW) and the take-up reel clockwise (CW).
- Brake – this signal stops or holds reel motion.

Two tape sensing devices are required per column. The top sensor looks for enough tape in the column, and the bottom sensor looks for too much tape. A simplified flow diagram of this operation is shown in figure 6-3. As can be seen from this diagram, if the top sensor is not covered, the tape is fed; if the bottom sensor is covered, the tape is taken up; if the top sensor is covered and the bottom sensor is not covered, the brake is applied.

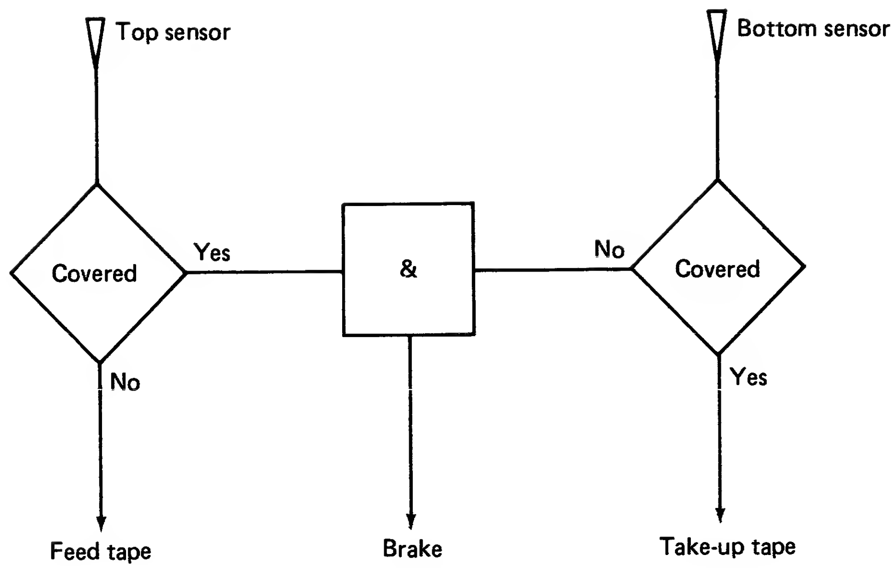

Figure 6-3. Sensor flow diagram

Mounting Tape

When tape is mounted on the tape transport, it is threaded from the right of the supply reel, around the supply idler roller, across the top of the supply vacuum storage column, between the head and pressure pad, across the top of the take-up vacuum storage column, around the take-up idler roller and onto the take-up reel from the left side (see figure 6-4). Several feet of tape must be wound onto the take-up reel. The tape is now ready to load.

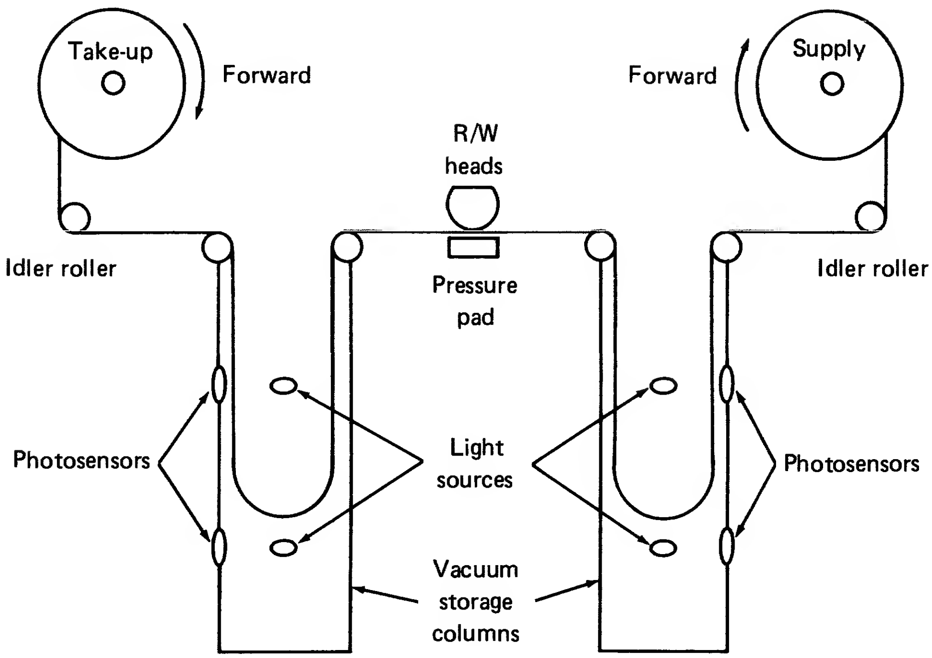

Figure 6-4. Tape path (simplified)

Load Switch

When the LOAD switch is pressed, the following actions occur:

- Vacuum motors come on and create a vacuum in the vacuum storage columns.
- The top photosensors in each vacuum storage column detect light; this causes the supply and take-up reel motor to turn and feed tape (the supply reel turns clockwise, the take-up reel turns counterclockwise).
- The combination of the atmospheric and vacuum pressure differences, discussed previously, and the feed tape motion of the reel motors causes the tape to load into the vacuum columns.
- As soon as the tape passes the top photosensors of each column, the reel motors are turned off or braked.
- The system now starts to look for the BOT (beginning-of-tape) marker.

Search for BOT

In the search for BOT, tape is removed from the supply vacuum storage column and placed in the take-up column. As the tape in the supply column reaches the top photosensor, the signal to feed tape is sent to the supply reel motor. This signal remains on until the tape covers the top sensor again. As more and more tape is fed into the take-up column, the tape will soon reach the bottom photosensor; when this sensor detects the tape, a signal is sent to the take-up reel motor to take up tape. This action continues until the BOT is found. The system is now ready to process data.

Read, Write, Search Forward

During a read, write, or search forward command, the tape will move from the supply reel to the take-up reel. The top sensor in the supply vacuum storage column initiates the feed tape signals. The bottom sensor in the take-up vacuum storage column controls the take-up reel to remove tape.

Read Reverse, Rewind, Unload

During the read reverse, rewind, or unload operation, tape is removed from the take-up vacuum storage column and placed in the supply column. Thus, the top sensor in the take-up column sends the feed tape signal to the take-up reel, and the bottom sensor in the supply column sends the take-up tape signal to the supply reel.

Tape Fault Sensors

There are two additional sensors or monitors in each column. These are tape fault sensors. One of these is above the feed tape sensor, the other is below the take-up tape sensor. The function of these sensors is to make sure the tape does not get too short or too long. If the tape fault sensors are triggered, the system will power down.

Maintenance

With an understanding of how the tape sensors operate, the following common sense maintenance procedures should be obvious.

- Vacuum columns should be kept clean to avoid false triggering of the sensors.
- The vacuum supply should be adjusted properly. If it is too light, tape will not be pulled into the column; if it is too heavy, too much tape will be pulled into the column.
- Proper adjustment and working condition of the sensors should be checked. False trigger action or no trigger action could result if the sensors are not working properly.

Summary

The correct and proper action of the feed tape and take-up tape sensors is essential to avoid damage to the tape and proper tape movement. If the feed tape sensor is faulty, the tape will be stretched or torn; if the take-up sensor is faulty, the tape will become creased—either may happen before the fault sensors have a chance to power the system down.

Tachometers

This activity introduces the tachometer and its function in the tape transport control.

Introduction

Tape speed control is an absolute necessity; if the speed of the tape changes, data will not be read or written in a uniform manner. The capstan drive controls the speed of the tape across the read/write heads. The speed of the tape coming off the feed reel, or going onto the take-up reel, also has to be controlled. If the speed of these reels is not controlled, the following problems occur:

- Too much tape is supplied to the feed loop box, causing a fault.
- Too much tape is removed from the take-up loop box, causing a fault.
- Tape is dragged across the heads, creating parity errors.
- Tape is stretched or torn.

Location of Tachometer

A tachometer is a device used to measure rotational speed; it is employed in some tape systems to slow down the reel motor when tape speed exceeds the capstan speed. The tachometer is usually located between the loop box and the reel (see figure 6-5).

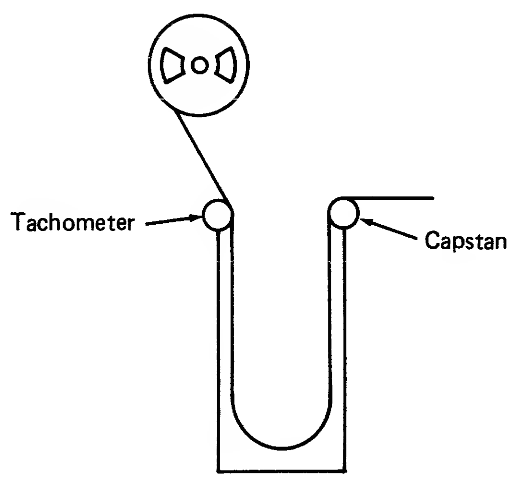

Figure 6-5. Typical tachometer location

Types of Tachometers

There are two types of tachometers used on tape transports. One type is an AC generator (see figure 6-6).

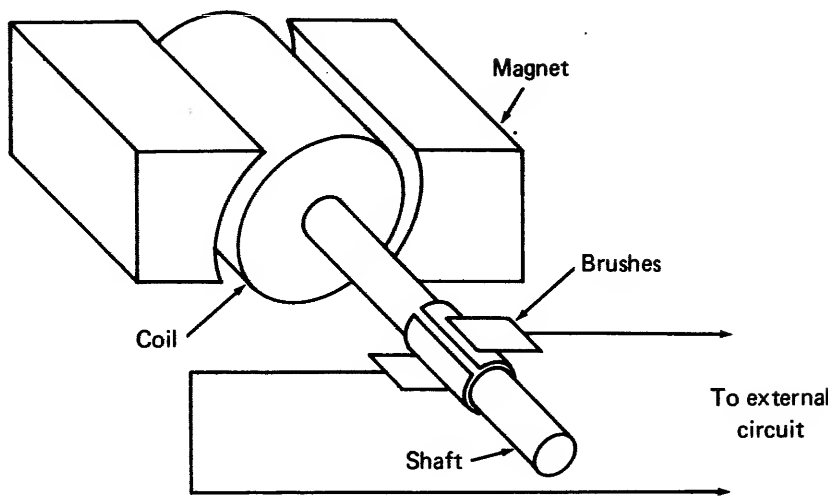

Figure 6-6. AC generator

According to basic AC theory, as the loop cuts the flux field in the magnets, a small voltage is induced in the coil. Leads connected to the brushes transfer this voltage to the external circuits.

This AC generator generates an output voltage and frequency directly proportional to the speed of the tape between the reel and the vacuum column. The faster the tape movement, the more voltage and higher frequency there is.

The other type of tachometer is the optical tachometer shown in figure 6-7. It is used more than the other type.

The optical tachometer uses a lamp, photodetector, and timing disk. The timing disk may be glass with the timing marks photoetched on the surface, or it may be plastic (or other material), in which case the marks are notches cut into the edge of the timing disk–similar to a gear. Care must be exercised when working with either of these disks to prevent damaging or scratching them.

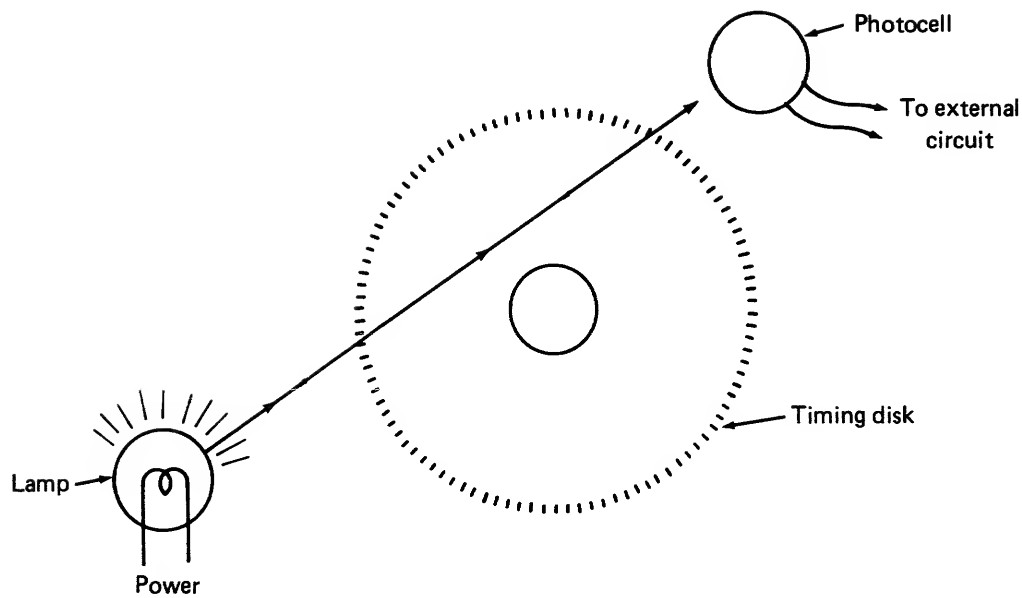

Figure 6-7. Optical tachometer

The optical tachometer generates output pulses for a light/no light condition. The frequency (repetition rate) of these pulses is directly proportional to the speed of the tape, which turns the timing disk.

Speed Control

Either type of tachometer controls the speed of the drive in a similar fashion. The output frequency is counted; if the counter is reset in a determined time frame, the speed of the tape is acceptable. If the counter reaches the maximum count without being reset, the speed of the tape is too fast. When this maximum count is reached, the logic will either send a brake command to the reel motors, or interrupt their applied voltage.

Rewind

During a rewind operation, the tachometer is disabled. This allows the tape drive to perform a high speed rewind. There is no data transfer occurring during a rewind operation; therefore, the speed of the reel motors does not have to be controlled.

Vacuum Sensors

This activity introduces vacuum sensors and their function in some tape transports.

Theory

At normal atmospheric pressure, the pressure on all sides of an object is equal. If this object is sealed, and a vacuum is "pulled" on the inside, the pressure on the outside of the object is greater than on the inside. (See figure 6-8.)

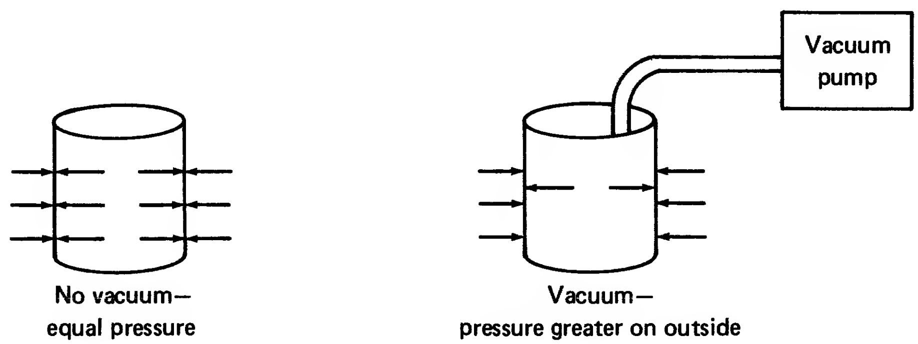

Figure 6-8. Pressure

This can be experienced by taking a vacuum sealed jar or can and opening it. The "woosh" sound you hear is the outside air rushing in; this is the pressure difference.

Differential Switches

Based upon the theory of unequal pressures, the differential switch was developed. Basically, this is a sealed container with a diaphragm and switch inside (see figure 6-9).

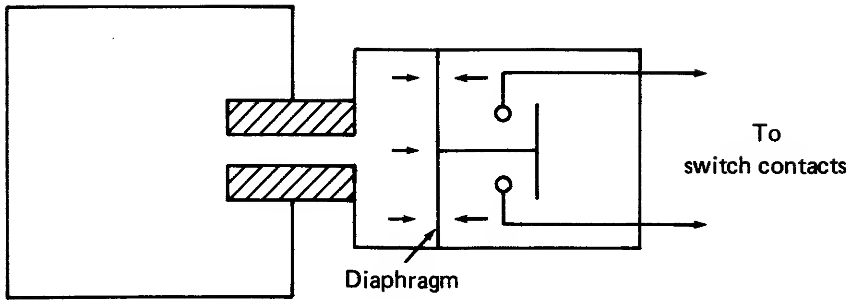

Figure 6-9. Pressure switch-normal pressure

Under a normal pressure condition, the pressure on both sides of the diaphragm is equal and the switch contacts are open.

If a vacuum is pulled on the pressure switch, as shown in figure 6-10, the diaphragm is "pushed" by the outside air, which causes the switch contacts to close.

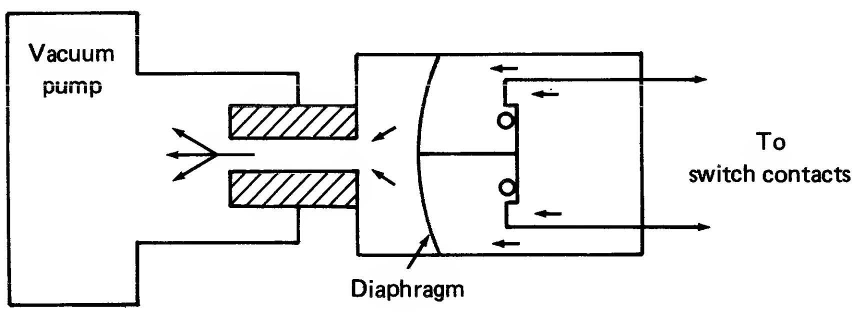

NOTE: Switch contacts may be reversed.

Figure 6-10. Pressure switch-vacuum supplied

Vacuum Sensor Function

The function of pressure switches in tape transports is to monitor the vacuum storage columns for tape. In figure 6-11, these switches are at the fault port locations.

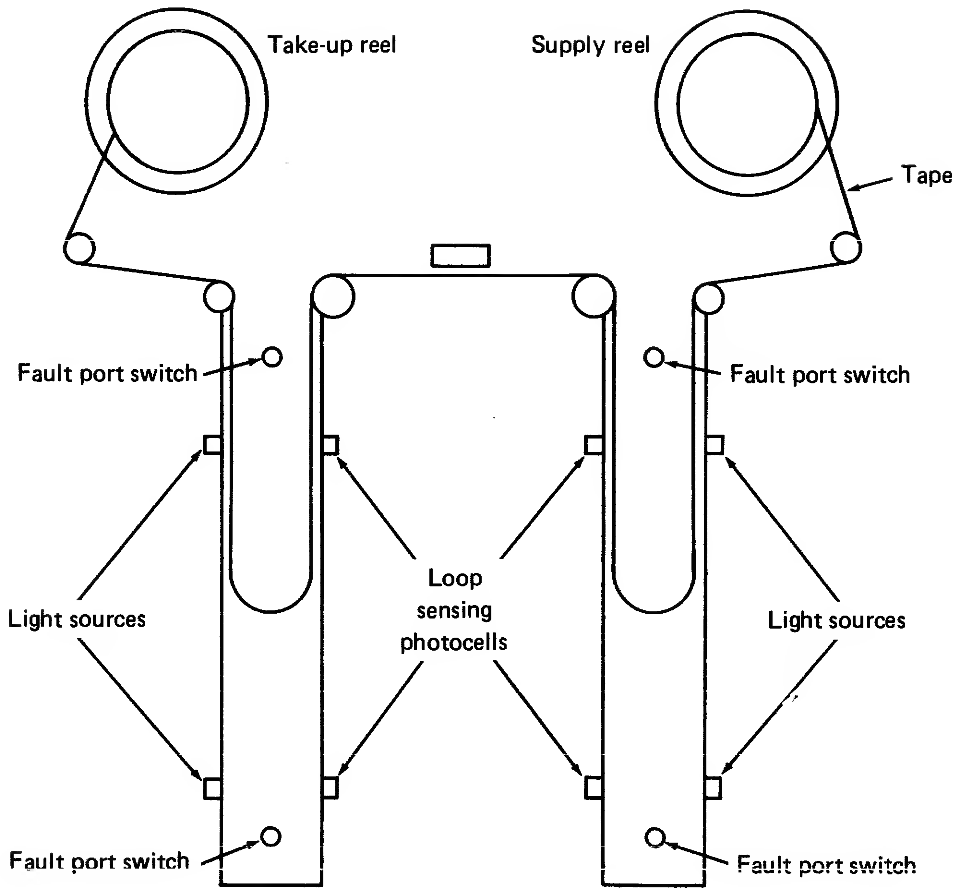

Figure 6-11. Pressure switch locations

When the LOAD switch is depressed, vacuum will be "pulled" on the vacuum storage columns. This will cause both fault switches in a column to close. In order to maintain a load/ready condition, the top fault switch must be open and the bottom switch closed.

The tape must load the columns quickly; the time is determined by the logic of the tape transport. This column load will cause the tape to go past the top fault port. This returns the atmospheric pressure to the pressure switch, causing it to open. If the tape remains between the two fault ports in each column, processing may continue. However, if the tape goes above the top port, or below the bottom port, a failure or fault will occur which will shut the system down.

Vacuum Sensor Operation

When tape is loaded properly, the normal operating tape length is between the two photosensors (see figure 6-12).

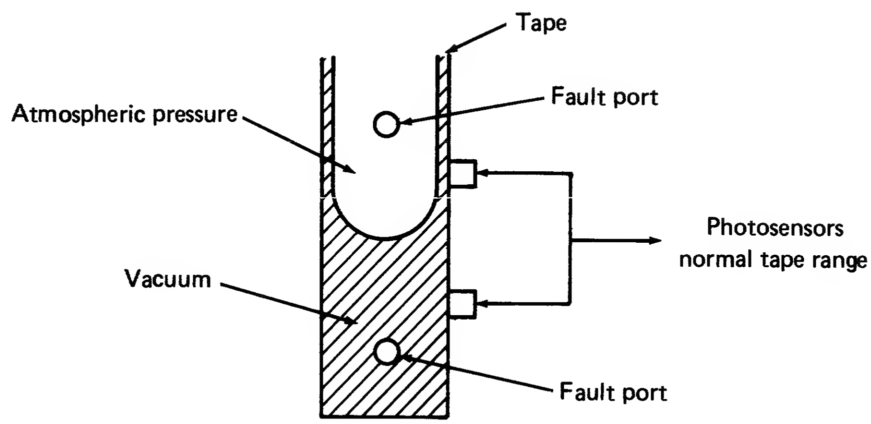

Figure 6-12. Normal tape range

Remember, the top photosensor will instruct the reel drive to add tape to the column and the bottom photosensor will instruct the reel drive to take up some tape from the column.

As you can see in figure 6-12, if a problem arises that prevents the reel drive from adding tape to the column, it is soon empty. This allows the tape to come to the top of the column, which, if not corrected, means the other reel has to take up the drag (weight) of the failing reel. This could result in stretched or broken tape. To prevent this, a fault port is installed (see figure 6-13).

As soon as the tape reaches the top fault port, a vacuum is applied to the diaphragm. This causes the switch to close, indicating a failure and shutting the system down to prevent damage to the tape.

Likewise, if the reel drive fails to remove tape from the column, tape piles up in the bottom. This creases or wrinkles the tape. Therefore, a bottom failure area is installed to prevent this. If the tape reaches this fault port, atmospheric pressure causes the switch to open, indicating a failure.

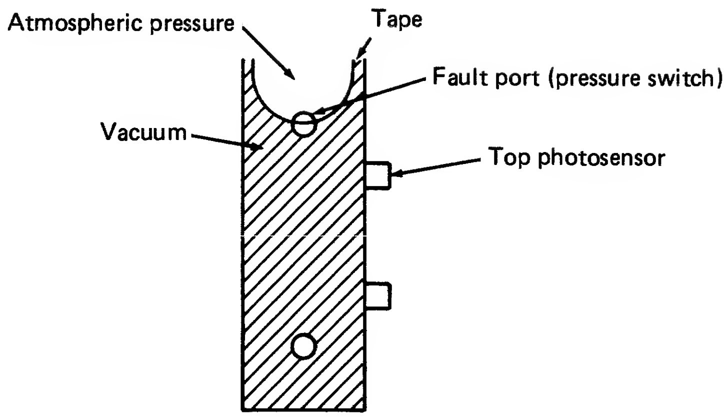

Figure 6-13. Top fault port

A simplified diagram of this circuit is shown in figure 6-14.

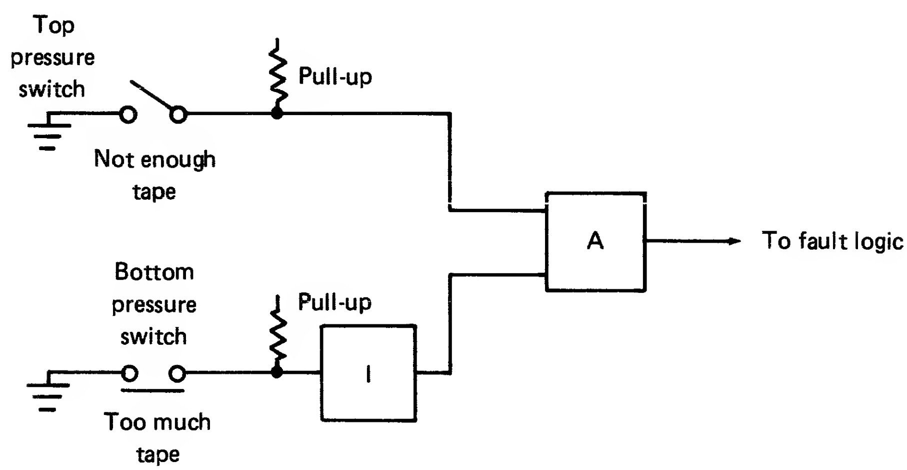

Figure 6-14. Fault circuit diagram

Photocells

This activity introduces photocells and light sources, uses in tape transports, and adjustment errors.

Photocells and Light Source

In general, the photocell is a silicon solar cell used to convert the radiant light energy into electric power. The output power of the cell is dependent on the radiant energy, or intensity, of the light source. Thus, the photocell acts as a switch that is activated by the presence or absence of light.

In electronic equipment, the light source is a common lamp. The light may be placed either in the direct path of the photocell or it may be reflected onto the photocell. Loop sensing in the vacuum storage column is an example of the direct path; an example of the reflected path is in the EOT/BOT sensors (see figure 6-15).

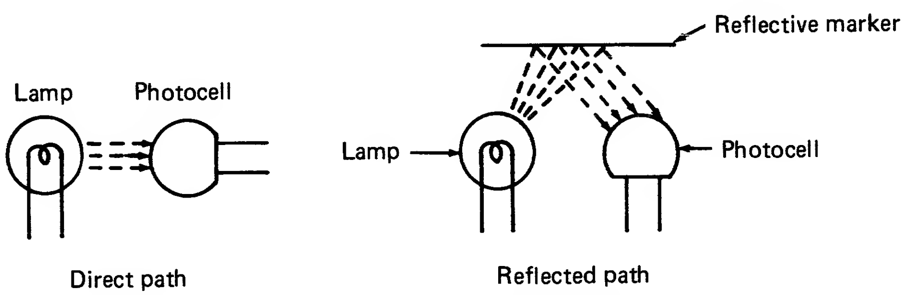

Figure 6-15. Lamp and photocell configurations

Another method of direct path light source is through the use of fiber optics. Fiber optics, also known as "light pipes," are made from a plastic-like flexible material which allows light to flow through them. The amount of light flowing depends on the diameter and purity of the light pipes. Any small fractures or cracks in the light pipes decreases the amount of light. Therefore, the ends of the light pipes are very smooth and buffed to a high polish. Care has to be exercised when handling light pipes; sharp angles will cause them to crack inside and reduce light flow.

Because fiber optics are flexible and transmit light, a single light source may be used, and it may be some distance away from the photocells (see figure 6-16).

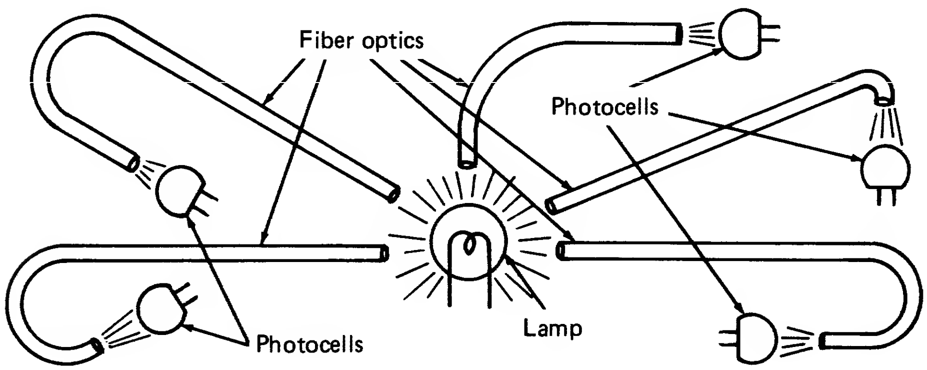

Figure 6-16. Fiber optics

Because the photocell converts radiant energy, it is not necessary to use only visible light, as from a flashlight or common lamp; invisible light, such as infrared, may also be used. The output of a photocell is usually connected to an amplifier. By using the infrared spectrum, it is possible to adjust the amplifier so that the visible light spectrum will have little or no effect on the photocell. It must be remembered, however, that the visible light spectrum contains some of the invisible spectrum; therefore, it is still possible to trigger the infrared configuration.

Uses in Tape Transports

Photocells have a wide use in tape transport systems. The following items utilize photosensors:

- EOT/BOT sensors
- Loop box sensors
- Tachometer

EOT/BOT Sensors

Located just before the read/write head are the EOT/BOT sensors. It is the function of these sensors to signal to the controller when the beginning or end of tape is reached. A highly reflective marker is attached to the tape, which, when positioned over the sensor, reflects light from the light source onto the sensor (see figure 6-17).

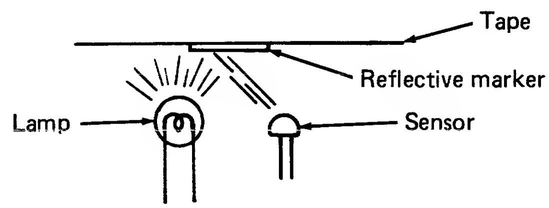

Figure 6-17. Tape sensor

There are two of these sensors, located side by side. The one toward the front of the transport is BOT, the other is EOT (see figure 6-18).

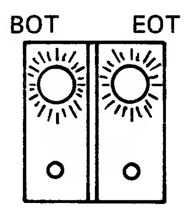

Figure 6-18. Sensor assembly

Loop Box Sensors

Located near the top and bottom of each vacuum storage column are the sensors that control the amount of tape in each column. The function of these sensors is to direct the reel drive motors to feed or take up tape.

Tachometer

The tachometer is located along the tape path. The function of this sensor is to measure the speed of the tape.

Adjustment Errors

Improper adjustment of the light sources or gains of the amplifiers could result in no operation or a false trigger condition.

If the light source is not aligned properly, or if the gain of the amplifier is too low, the signal would not change states with a light/no light condition.

If the amplifier gain is too high, the signal could oscillate or trigger due to extraneous light, thus causing timing or false start conditions.

Block 7

Tape Sensors and Read/Write Heads

Reel Brakes

This activity introduces some aspects of two types of reel brakes used on magnetic tape transport reel drive motors. The two types of reel brakes are magnetic particle brakes and single disk brakes.

Reel Brakes

Reel brakes are found on the rear of the reel drive motors. The functions of the reel brakes are: 1) to stop the reels after a FEED or TAKE-UP signal; this prevents excessive tape from being loaded into the vacuum storage columns, and 2) to hold the tape reels stationary; this prevents tape from creeping across the read/write heads.

Magnetic Particle Brakes

A unique type of reel brake is the magnetic particle brake. Magnetic particle brakes are a tightly sealed unit containing very fine powder. This powder, when magnetized, packs together to form a solid block. The inside of the brakes is designed so that when the powder magnetizes, it inhibits movement of the shaft that is part of the motor. The unmagnetized powder has little effect on the turning motor shaft.

Single Disk Brakes

Magnetic particle brakes have a tendency to leak the powder, so another design has been developed and is used on later transport models. Referred to as single disk brakes, these brakes consist of a brake coil that bolts onto the end of the motor and a brake armature that connects to the motor shaft. These two brake parts are very close together when adjusted properly. To stop the reel motor, current is passed through the brake coil. This current produces a magnetic field that pulls the brake armature to the brake coil face. The facing on the brake coil is a material that creates friction when it contacts the rotating brake armature. Thus, it stops the motor shaft.

Maintenance Procedures

Maintenance procedures for reel brakes consist of maintaining the proper gap between the brake armature and the brake coil, and insuring that the amount of current to the coil is sufficient to produce 60 ($\pm$ 5) inch-pounds of torque. The torque should be measured after the transport has been running continuously for 30 minutes. A variable resistor housed in the reel drive power supply is adjusted to obtain the 60 inch-pounds of torque. A special adapter is needed and a 0-100 inch-pounds torque wrench is used to measure the torque applied to the reel motor. After a new brake is installed, this torque adjustment should be made after 8, 16, 40, and 80 hours of operation.

In the latest tape transport models, reel drive motors have no brakes. Instead, the motor is designed to allow a simple reversal of current to stop the reel motion. This type of "braking" requires sensing electronics to insure that the right amount of current is produced for the right amount of time.

EOT/BOT Sensors

This activity describes the position and operation of the photosensors and the EOT/BOT sensors.

Position of Photosensors

Sensors for EOT (end-of-tape) and BOT (beginning-of-tape) utilize nearly the same photocell/light combination used for loop sensing in vacuum columns. EOT/BOT sensors, however, are mounted side by side, under the tape path and just beside the tape read/write heads. The photocell and light are housed in the assembly, both facing the top glass-covered center of the assembly. (See figure 7-1.)

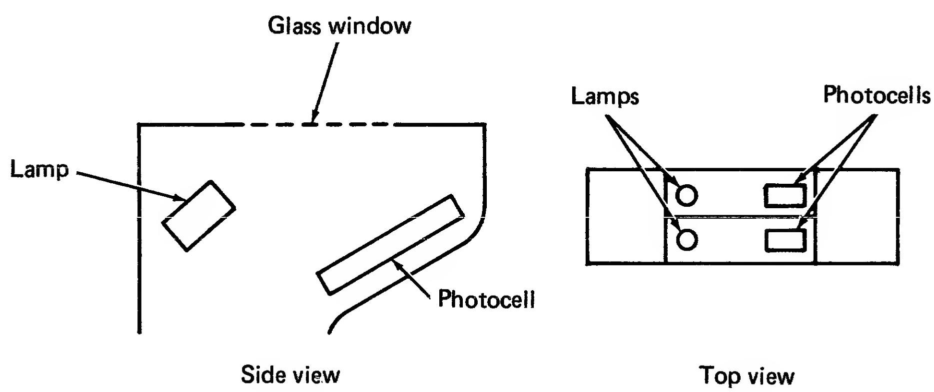

Figure 7-1. EOT/BOT assembly

Operation of Sensors

If reflective material is placed over the glass window on the top of the assembly, light will be reflected to the photocell, making its associated amplifier output a 1. (See figure 7-2.)

EOT and BOT sensors send signals to EOT and BOT circuitry. These circuits contain photocell amplifiers to change current of the photocell to logic signals, flip-flops to store the fact that the EOT or BOT has been encountered, transmitter circuits to send the signal to the tape controller, and relays to light indicators on the tape transport operator and maintenance panels.

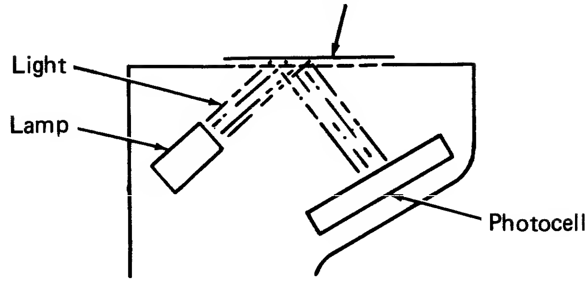

Figure 7-2. Light reflected onto photocell

Position of BOT/EOT Markers

The "beginning-of-tape" (also known as load point), and "end-of-tape" markers are placed near the beginning and the end of the tape to enable sensing of the usable portion of the tape by the photocells. Adhesive on one side of the markers secures the markers to the tape. The other side of the markers is coated with aluminum that is highly reflective. The markers are approximately 1.2 inches long and 0.2 inch wide and are placed on the uncoated side of the tape (underside of the tape when mounted on the tape deck). The end-of-tape marker is placed on the edge of the tape nearest the tape deck; the beginning-of-tape marker is placed on the outer edge of the tape. Figure 7-3 shows the precise location of the markers on the tape.

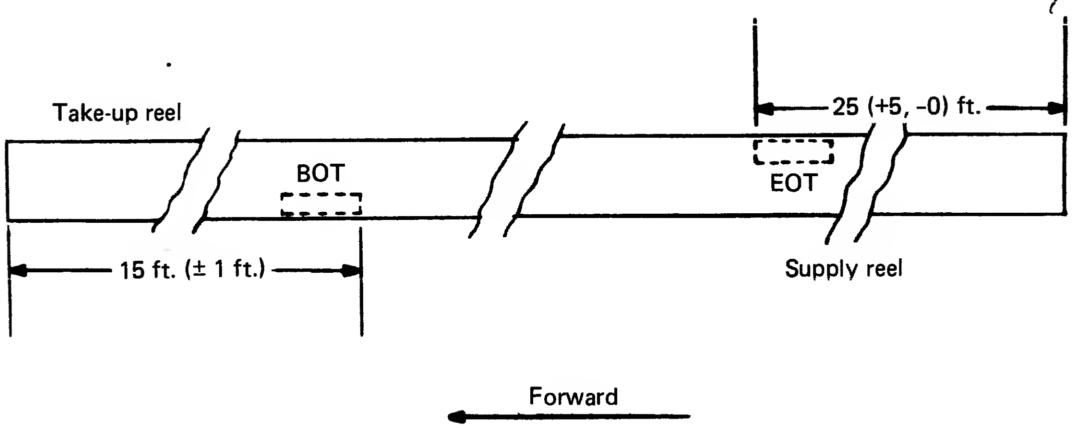

Figure 7-3. Reflective marker position on tape

Read/Write/Erase Heads

This activity discusses the location and the mechanical and electrical alignment of the read/write/erase head assembly used on magnetic tape transports. Variations may exist in transports but they are minor.

Location of Read/Write/Erase Heads

The read/write head is mounted with the erase head and one or two tape cleaners on the head assembly between the forward and reverse capstans. This head assembly is mounted on a special head mounting plate that allows vertical and horizontal alignment. Figure 7-4 shows the vertical relationship required of the various parts of the head assembly.

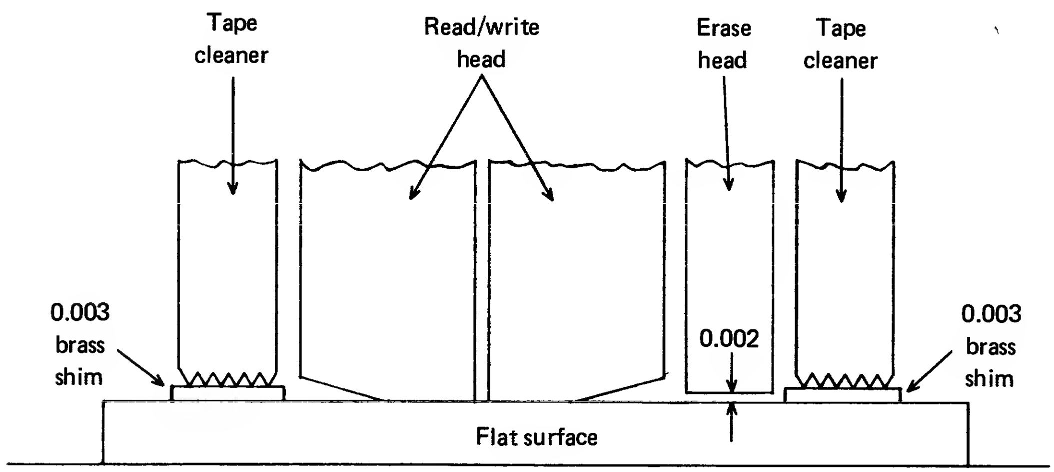

Figure 7-4. Head assembly components

The vertical position of the read/write head is critical and requires special equipment to insure correct height from the tape. This distance is necessary because of the magnetic lines of force required for read/write operations. Placement of the head assembly along the tape path is not critical; however, the data written or read must not be skewed. In other words, the data on one side of the tape must not be early while data on the other side of the tape is late. (See figure 7-5.)

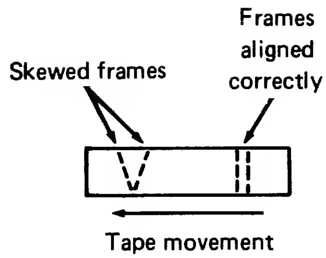

Figure 7-5. Tape with frames showing

As was shown in figure 7-4, tape cleaners are mounted on the tape head assembly. These cleaners are connected to vacuum lines that remove dust and other particles from the tape before the tape gets to the read/write and erase heads.

Just below the head assembly is a movable pressure pad which holds the tape against the read/write head with air pressure. During rewinding, the pressure pad is retracted from the head assembly.

Mechanical Alignment

The head mounting plate is beveled toward each side to allow proper head alignment. (See figure 7-6.)

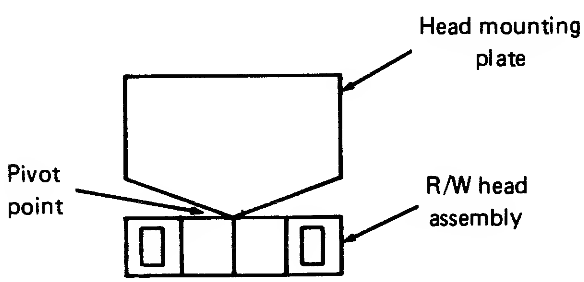

Figure 7-6. R/W head and head mounting plate (view from top)

An adjustment screw will rotate the head on the pivot point, allowing the head to be adjusted so that it's parallel to the tape, thus preventing the skewed data.

Electrical Alignment

To insure proper frame alignment, a special master alignment tape is read that allows a person to scope the preamplifiers and compare the time each track of data is encountered. Also, these same preamplifiers can be adjusted to maintain uniform amplification of all tracks.

Maintenance Concepts

As with any electrical and/or electronic piece of equipment, dust and dirt are enemies. However, with magnetic storage equipment, these impurities can be especially harmful. A speck of dust can easily create parity errors. This activity introduces some of the maintenance procedures necessary to keep magnetic tape equipment operational, and lists specific dos and don'ts concerning the handling of magnetic tape.

Preventive Maintenance Procedures

Operations personnel are responsible for the basic cleanliness of the equipment. Vacuum columns, rollers and the BOT/EOT sensors should be kept clean; in other words, the entire tape path should be kept free of dust and dirt. Particular attention should be paid to the read/write head to keep parity errors to a minimum. The preventive maintenance usually performed includes:

- Weekly inspection of vacuum and pressure gauges and capstan drag
- Biweekly electrical start/stop timing, photocell inspection
- Monthly filter cleaning/replacement and inspection of the reel drives
- Quarterly inspection of tape cleaners, voltage levels, plastic tubing, pump speed, and tape path alignment

Each time preventive maintenance is performed, run the diagnostic tests to verify that the equipment is still operating at peak performance.

Handling Magnetic Tapes

Certain precautions should be taken when handling magnetic tape. The following list of dos and don'ts will be helpful to you when you handle magnetic tape:

DO	DON'T
DO keep reels and containers clean.	DON'T allow any portion of tape to contact floor.
DO use care in placing reflective markers on tape.	DON'T lay tape on dirty surface when applying reflective markers.
DO start tape flat on take-up reel.	DON'T allow tape to fold or protrude through reel flange cut-out when starting on take-up reel.
DO pick up tape reels by holding reel at hub and outer edge of reel flange.	DON'T squeeze tape reel flanges when picking up or holding reel.
DO avoid contact with tape through cut-out areas of reel flanges.	DON'T nick or damage tape edges at reel flange cut-outs.
DO inspect tape and reel closely if dropped.	DON'T handle tape carelessly.
	DON'T unload and handle tape when not completely wound on reel.
DO seat reel squarely on transport hub when mounting.	DON'T mount reel on transport hub in cocked manner.
DO keep transport clean, especially heads and tape transport mechanism.	
DO keep tape in a clean container when not on transport.	DON'T allow tape to lie uncovered or exposed to contamination.
DO use labels for identification of reels which will not leave residue when removed.	DON'T use eraser in changing identification written on reel labels.

Causes of Problems

Failure to handle tape properly or failure to apply preventive maintenance procedures could result in unscheduled maintenance calls. Following is a list of tape problems and possible causes:

TAPE CONDITION	POSSIBLE CAUSE
Creases or indentions in tape surfaces.	Wrinkled reflective markers. Tape not started flat on take-up reel hub. Improper handling of tape. Dirty take-up reel hub. Tape end allowed to drop on floor.
Creased, nicked, or damaged tape edges.	Reel held by flanges when mounting or removing from transport; flanges squeezed or pressed against tape edges. Fingers nicking or pressing against tape in cut-out areas of reel. Warped reel. Tape reel has been dropped. Tape unloaded and handled when tape is not completely rewound to supply reel or take-up reel. Reel not seated properly on transport hub. Transport not adjusted properly.
Accumulation of foreign material on tape surfaces.	Tape transport not being cleaned properly. Tape not being replaced in closed containers when not in use. Dust or dirt in tape containers. Residue left on reels when labels are removed. Eraser used in changing information contained on reel labels. Fingerprints on tape surfaces.

TAPE CONDITION	POSSIBLE CAUSE
Accumulation of foreign material on tape surfaces (continued).	Smoking or eating in machine room.
	Unnecessary traffic in machine room.
	Improper filtering of machine room air.
	Improper or insufficient cleaning of machine room.
Wavy tape edge or stretched appearance along edges.	Reel hubs or idler wheels improperly adjusted.
	Reel not mounted squarely on transport hubs.
	Warped reel.
	Improper insertion of file protect ring, preventing proper reel mounting.
	Improper handling of tape when mounting or removing reel from transport.
Loss of recorded data.	Accumulation of foreign material on tape surface.
	Creased or damaged tape.
	Nicked tape edges.
	Wavy or stretched edges.
	Accidental erasure.
	Dirty read/write head.
	Skewed heads.
Failure to load tape.	Foreign material in vacuum columns.
	Dirty loop box sensors.
	BOT/EOT reflective markers incorrectly positioned or missing.
	Dirty BOT/EOT sensors.
	Improper vacuum.
Failure to write correct data.	Skewed tape.
	Skewed tape heads.
	Dirty tape heads.
	Creased or damaged tape.

Specific maintenance procedures are described in each tape transport manual and should be strictly followed.

Documentation

Every piece of equipment should have up-to-date maintenance logs. These logs let both the customer and maintenance personnel know when equipment maintenance is due and completed. They assist in the detection of potential problems by letting you know when certain adjustments become too frequent, and they give you a general idea about how the equipment is performing.

When equipment has unscheduled maintenance performed on it, another item of documentation used is an equipment operations report. This report is usually filled out by the operations personnel. When a problem arises, maintenance personnel correct the problem and identify the corrective procedure on the report. By monitoring these reports, it is possible for personnel to detect and correct a malfunction or make a necessary design change before drastic failures occur.

Block 8

Introduction to Disk Storage Devices

Introduction to Disk Storage

Working with disk storage devices requires that you understand technical terms used by other persons in the field. This activity defines common terms associated with disk storage devices.

Rotating Magnetic Disk Storage

This is the general term for all types of disk storage devices. Rotating magnetic disk storage includes all classes of computer memory based on rotating magnetic disks. This type of memory is generally classified as random access memory, because any block of data can be accessed directly, or non-sequentially. However, data within the block is read or written sequentially and cannot be randomly accessed. Data is organized, in ascending order, in bits, bytes, sectors, tracks, and cylinders.

Bit

A bit is the simplest, or lowest order, unit of logical or digital information. A bit is either 1 or 0. On disk storage, one flux change in the recording media does not always represent just 1 bit.

Byte

On disk storage, a byte is a specific number of consecutive bits, usually 8. A byte represents a unit of binary coded information, such as a single character, or number.

Track

When a head is held in a stationary position over a rotating disk, it creates a circular path. The string of bytes that can be recorded on this circular path is called a track.

Sector

A sector is a subdivision of one track, although it could fill up one track. It is usually only a portion of the arc of the circular track and it has an identifiable beginning and end. It is the smallest block of data that can be randomly accessed on disk storage. All the bytes within a sector must be read sequentially.

Disk Pack

Many disk storage devices use several disks mounted on a spindle and stacked one on top of the other. This array of disks is called the disk pack.

Cylinder

On disk storage devices that use several disks, or a disk pack, each disk surface has at least one head. All heads move together, in and out, over the disk pack. When one head is positioned over one track on a disk, the other heads are also positioned over tracks on other disks. A cylinder is the group of tracks that can be accessed when the heads are held in one position.

Format

When signals are written on and read from the disk within a sector, they must be written and read in a specific, standard pattern that the computer can utilize. This standard pattern is the format. The format of a sector divides it into an address portion and a data portion.

Address

The address portion of a sector gives it a unique identity. It generally includes the cylinder (which locates the group of tracks the data is in), the head (which identifies a particular track in the cylinder), and the sector (which pinpoints the location of the data on the track). The manner in which these signals are recorded is defined by the format.

Data

The data is the usable information stored in the sector. It is stored in byte units.

Address Mark

Just before the address signals, a portion of the track may be erased. This address mark identifies the beginning of an address field.

Index

The index is a signal that identifies a unique position of the disk; it marks the beginning of a revolution. The index signal is generated either by a magnetic pulse recorded on the disk or by a notch cut into the edge of the disk.

Sector Mark

The sector mark is a signal that identifies the beginning of a sector. It can be generated either by a notch cut into the edge of a disk or by a magnetic pulse recorded on the disk.

End-of-Travel (EOT)

The band of tracks holding data does not cover the entire surface of the disk. There is an area on the outer and inner edges of the disk that does not hold data. These two bands mark the end-of-travel, abbreviated as EOT, for the heads. The inner area is called the FORWARD EOT, because it is the end of forward head motion, and the outer area is called the REVERSE EOT, because it is the end of reverse head motion.

Return to Zero (RTZ)

The cylinders of a disk pack are numbered consecutively starting with 0 as the outermost cylinder. Return to zero, abbreviated as RTZ, is an automatic function of the disk storage device that moves the heads from wherever they are on the disk back to cylinder 0.

Retract

On some disk storage devices the disk pack can be removed from the deck and be replaced with another pack. On these devices, the heads must be moved completely off the disk pack before the pack can be removed or a new one installed. The retract position is the position the heads are in when the disk pack can be moved.

Seek

On disk storage devices with moveable heads, the heads move from cylinder to cylinder to locate an address and write or read data. The operation of moving from one cylinder to another is called a seek operation because the device starts with the heads on one cylinder, then gets instructions to look for, or seek, the next cylinder.

Seek Error

During a seek operation, an error might occur. As a result, the heads wind up on the wrong cylinder or the device cannot find the specified cylinder. When this happens, the device indicates to its controller that a seek error has occurred. The controller responds by telling the device to return to zero and try again.

Spare Tracks

Disk storage devices do not usually use all the tracks available on a disk to store data; they keep some tracks reserved as spare tracks. A primary data track may be defective because of damage to the disk recording surface, or the customer engineer may want to perform tests on the disk storage device. A spare track can be used to hold the data from the defective track, or it can be used for tests by the customer engineer. The inner tracks of a disk are usually used as spare tracks.

Disk Storage Classifications

There is a wide array of disk storage devices now being manufactured and marketed, but they can be classified in several ways. This activity lists the means of classifying disk storage devices and the categories within those classes. This activity also lists combinations of characteristics that represent disk storage devices available in 1978.

Types of Disks

Disk storage devices use disks that are either rigid or flexible.

Rigid Disk

Rigid disks are usually made of aluminum with a coating of iron oxide painted on. The vast majority of disks in use today are 14 inches in diameter and they are mounted in a disk pack of anywhere from 1 to 20 or more disks. The first disk storage devices used rigid disks and they are still the most popular and versatile.

Flexible Disk

For certain applications, the flexible disk, usually known as a "floppy," is the most appropriate device. Flexible disks are made of 1.5 mil Mylar coated with iron oxide. They are essentially disk-shaped pieces of computer magnetic tape 5-1/4" or 8" in diameter. Floppy disks are contained in a slip-cover jacket like a record jacket; however, the floppy disk is not removed from its jacket when it is used. Flexible disks are less expensive and easier to handle than rigid disk packs, but they hold less data and are slower than rigid disks.

Types of Head Configurations

Disk storage devices use heads that are either movable or nonmovable.

Movable Heads

Devices with movable heads generally have one head for each disk surface. A positioner device, also called an actuator, moves the head from track to track over the disk. Some moveable head devices have two heads for each disk surface, which cuts in half the distance each head has to move.

Nonmovable Heads

Disk storage devices with nonmovable heads are also called "head-per-track" machines because they have one head for each track on a disk. Because the heads are relatively large, this severely limits the number of tracks the device can record on a disk, and, thus limits the amount of data that can be recorded. A movable head device can record many more tracks on a disk of similar size. However, the actuator and its associated electronics, which are required to move the heads, are complex and expensive. A head-per-track machine does not require an actuator. Also, mechanically moving a head from track to track is slower than electronically switching from head to head, so head-per-track devices are faster.

Types of Disk Packs

Disk storage devices use disk packs that are either removable or fixed.

Removable Disk Pack

Devices using removable disk packs enable the customer/operator to take out the disk pack and install a different one. In this way, a library of data on disk packs can be maintained, or a pack can be moved from machine to machine. When removed, the pack is held in a protective canister or is mounted in a permanent cartridge. Removable packs are used only on devices that have movable heads because the heads must be retracted to allow removal of the disk pack.

Fixed Disk Pack

Disk storage devices with fixed packs do not allow the customer/operator to change disk packs. However, a customer engineer can replace defective or malfunctioning fixed disks. On a fixed pack device, disks, spindle, heads, deck, and sometimes the actuator are all mounted in a sealed module. The sealed module prevents environmental contaminants, such as dust particles or moisture, from interfering with operation of the disk and heads. A fixed pack device can use either movable or non-movable heads or a combination of both.

Capacity

Capacity refers to the number of bytes the device can store. Small capacity devices store under 10 million bytes, or 10 Megabytes (10 MB). Medium capacity devices store 10 MB to 100 MB per disk pack. Large capacity devices are capable of storing over 100 MB per pack. Capacity is one of the primary characteristics of a disk storage device a customer considers when deciding which class of disk storage device to use.

Combinations

The preceding categories (types of disks, head configurations, disk packs, and capacity) can be combined in many ways to describe disk storage devices. Some combinations simply aren't produced because they aren't practical or fill no need. Other combinations are available in 1978; these are described below.

Small Capacity (Less than 10 MB Per Pack)

Flexible Disk, Movable Head, Removable Pack

Although this device is slow and has a low capacity for data, it is growing in popularity because it is inexpensive and useful for storing programs to initiate computer systems.

Rigid Disk, Nonmovable Heads, Fixed Pack

Although this device has a low capacity for data, its access time is fast. It is useful in real-time applications where rapid data transfer times are a necessity.

Rigid Disk, Movable Head, Fixed Pack

This device is useful as main memory for small computer systems, bulk memory, or for ancillary data storage such as table lookups.

Rigid Disk, Movable and Nonmovable Heads, Fixed Pack

This device combines advantages of both movable and nonmovable head devices. It generally has several disks with one disk accessed by nonmovable heads and the others accessed by movable heads. It combines rapid access to some data with large storage capacity on the other disks.

Rigid Disk, Movable Heads, Removable Pack

This versatile device allows building up a library of disk packs or moving a disk pack from location to location.

Medium Capacity (10 to 100 MB Per Pack)

Rigid Disk, Movable Heads, Fixed Pack

This device is similar to the same combination under small capacity. The fixed pack is generally sealed, preventing environmental contaminants from interfering with operation of the disk. This allows heads to be closer to the disk, which enables the device to store more data on a disk.

Rigid Disk, Both Movable and Nonmovable Heads, Fixed Pack

This device has the same uses and advantages as the same combination in small capacity, except it has a larger data storage capacity.

Rigid Disk, Movable Heads, Removable Pack

This is a common device; it is a standard workhorse for computer systems.

Large Capacity (Over 100 MB Per Pack)

The same combinations available in medium capacity devices are available in large capacity devices. The larger capacity is needed for larger computer systems or for archival storage of data.

Common Performance Parameters

This activity defines the most common performance parameters used to describe disk storage devices. This activity also lists the particular disk storage devices that have become industry standards and gives their performance parameters. By reading this activity and studying its contents you should be able to list and define common performance parameters. You should also be able to recognize the historical importance of particular disk storage devices when they are mentioned.

Performance Parameters

Disk storage devices have performance parameters that describe to a customer the capabilities of a particular device and enable the customer to compare one device with another.

Latency

Latency is the time required for a specific sector of data to rotate into position under the head once the heads are on the correct cylinder. It is usually given as the average latency, which is the amount of time required for the disk to move one half rotation once it is up to speed. Latency is one indication of how quick the device is.

Seek Time

Seek time is the time required for the actuator to move the heads from their present position to the next desired cylinder. It can be given as the minimum seek time, which is the time required to move the heads one track. Or, it can be given as the maximum seek time, which is the time required to move the heads from the outermost track to the innermost track. Usually, it is also given as average seek time, which is the average of all possible seek times.

Access Time

Access time is the sum of the average latency and average seek times. It is a general indication of the speed of the particular device.

Tracks-Per-Inch (TPI)

Tracks-per-inch is an indication of how densely data can be packed on a disk. Given a radius on the storage disk, TPI tells how many tracks cross a one-inch segment of the radius.

Bits-Per-Inch (BPI)

Given any track on a disk, bits-per-inch indicates how many bits can be recorded within a one-inch segment of that track. Data is recorded at the greatest density on the innermost track of the disk. The BPI figure for any machine is the density of data on the innermost track.

Transfer Rate

The transfer rate is the rate at which data can be read from or written onto disk storage. It is measured in bytes/second or bits/second. Transfer rate is an important indication of the speed of a disk storage device.

Capacity

The capacity of a disk storage device is the total number of information bits, or bytes, that can be stored. Capacity is usually specified as formatted data bytes per spindle.

Read Error Rate

Given an amount of data, the read error rate indicates the maximum number of bits that can be read erroneously. A read error can be caused by a tiny fault on the disk, dust, or an electronic malfunction. The read error rate is usually extremely small, such as 10-9, which means that only 1 bit in every billion can be read erroneously. There are also two types of read errors. One type of read error, called a soft error, can be corrected by simply re-reading the data. The other type of error is a hard error, which cannot be corrected by re-reading the data. For every 100 soft errors, there will usually be one hard error. Some typical soft and hard error rates are:

	Soft	Hard
Flexible disk	10-9	—
Rigid disk, moveable heads	10-10	10-12
Rigid disk, nonmovable heads	10-11	10-13

Errors up to a certain size can be corrected without re-reading if an error correction code is utilized when recording the data.

Cost/Bit

Cost-per-bit is the total cost of the disk storage divided by its capacity. It can be given as either rental cost/bit or purchase cost/bit to reflect the difference between leasing or buying equipment. It is an important parameter to the customer because it enables the customer to compare the device to other disk storage devices and to other memory devices.

Industry Standard Units

Often when persons in the peripheral devices field discuss disk storage devices, they mention certain devices as standards, such as the IBM 2314, IBM 3330, Control Data 9760, or IBM "Winchester." Each of these machines represents a new generation or marked advance in disk storage technology which other manufacturers then proceeded to copy and use. Table 8-1 lists industry standard machines in terms of the performance parameters defined above.

Look at the IBM 2314 specifications. The 2314 used a hydraulic device to position the heads and represented a major advance in disk storage technology. The transfer rate is more than double that of the previous device and its capacity is more than four times that of the 2311.

The next major advance was the IBM 3330 device. It used a magnetic system to position the heads, which enabled TPI to almost double. Transfer rate almost tripled and capacity again increased almost fourfold. The 3330 was introduced in 1970, only five years after the 2314, and represents a technology still commonly used today.

Three years later, the IBM 3340 was introduced; it had both movable and nonmovable heads and had two movable heads per disk. Average seek time, TPI, BPI, and transfer rate were all improved. However, its capacity was smaller, reflecting the need for disk storage devices to use with minicomputers and small computer systems. The technology of putting two heads per disk surface within a sealed cartridge is usually referred to as 'Winchester' technology.

In 1974, the Control Data 9760, referred to as the "Storage Module Drive" or SMD, became an industry standard for minicomputer uses. Improved head technology give it a high BPI and very high transfer rate.

TABLE 7-1
History of Industry Standard Rigid Disk Storage Units

Type	IBM 350	IBM 1311	IBM 2311	IBM 2314/2319	IBM 3330	IBM 3330-11	CDC 9760/9762 (SMD)	IBM 3340 (Winchester)	IBM 3350
Average Latency	25MS	20MS	125MS	12.5MS	8.3MS	8.3MS	8.3MS	10MS	8.3MS
Average Seek Time	600MS	150MS	75MS	60MS	30MS	30MS	30MS	25MS	25MS
TPI	20	50	100	100	192	370	192/384	300	476
BPI	100	1020	1110	2200	4040	4040	6038	5600	6425
Transfer Rate	10KB/S	78KB/S	156KB/S	312KB/S	806KB/S	806KB/S	1200KB/S	885KB/S	1198KB/S
Capacity/Spindle	5MB	2MB	7.25MB	29.2MB	100MB	200MB	40/80MB	35/70MB	317.5MB
Year First Shipped	1956	1962	1964	1965	1970	1974	1974	1973	1976
Configuration	FP/MH	RP/MH	RP/MH	RP/MH	RP/MH	RP/MH	RP/MH	RP/MH&NMH	FP/MH&NMH

MH – Movable Head

NMH – Nonmovable Head

RP – Removable Pack

FP – Fixed Pack

Summary

The performance of disk storage devices can be described with several common performance parameters. Certain disk storage devices represent advances in technology and become industry standards.

Common Assemblies

Although each model of disk storage devices is somewhat different, all disk storage devices have common types of assemblies. This activity lists and describes the assemblies that are common to all disk storage devices. By reading this activity and studying its contents you should be able to give a functional description for any of the common assemblies listed below.

Assemblies Common to Most Disk Storage Devices

Refer to figure 8-1 as you read the following assembly descriptions.

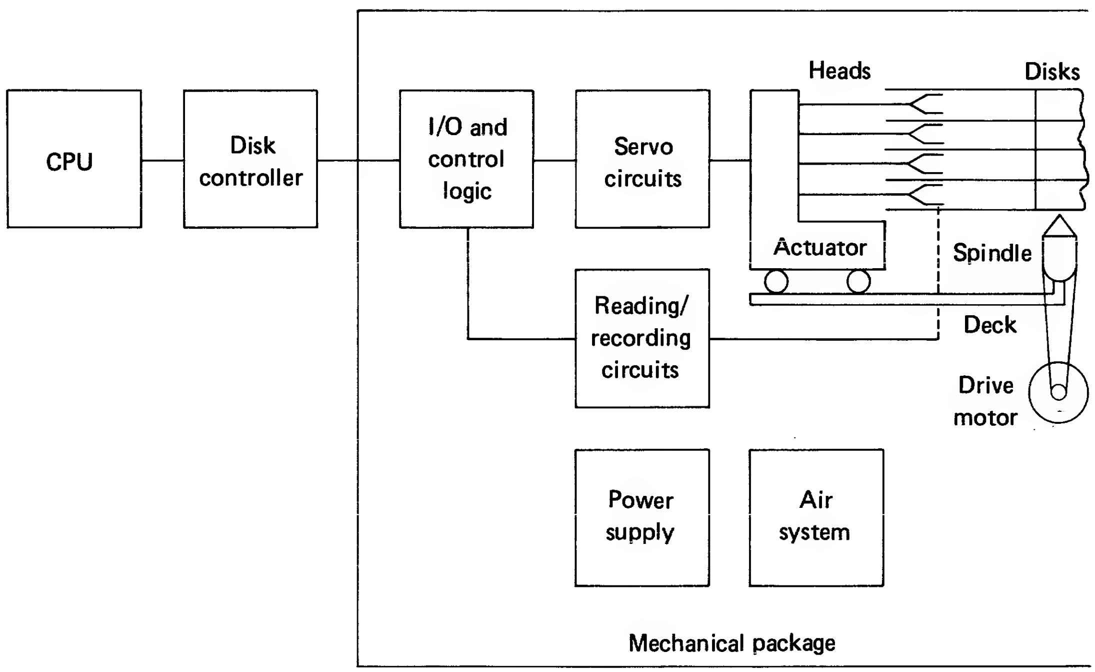

Figure 8-1. Typical disk drive configuration

Disks

The disks are the data storage medium. Made of aluminum, they are coated with iron oxide and stacked in a disk pack that can have from 1 to 20 or more disks.

Heads

The heads are magnetic transducers that read data from and write data onto the disks. On most disk storage devices, the heads "fly" over the disk surface, meaning that a thin film of air acts as a bearing between the disk surface and the head. On some floppy disk devices, the head actually touches the disk.

Spindle

The spindle is the rotating assembly which holds the disks in place in the disk pack.

Actuator

Only disk storage devices with movable heads have actuators. The actuator is the assembly that holds the heads, generally in a 'comb' arrangement, and moves them in and out over the disk surfaces.

Deck

The deck is the framework that holds both the disk and spindle assembly and the actuator.

Drive Motor

The drive motor is an alternating current motor that spins the disks on the spindle. It is usually connected with a belt.

Air System

The air system blows filtered air into the disk pack area of the device to maintain positive air pressure. This keeps dust particles and other contaminants out of the disk pack area. The space between the heads and disks is extremely small, and any particles could cause damage to the disk or heads. The air system may also cool the electronics. Generally, only rigid disk devices have air systems.

Electronics

The disk storage device has electronic circuits for reading and writing, disk control, and input/output logic, and devices with movable heads have circuits for servo control of the actuator.

Cables

Cables connect the disk storage device to its controller. They transmit both data and control signals; several disk units may be connected to a single controller or a disk unit may be accessed by more than one system.

Power Supply

The power supply provides direct current to the drive electronics and alternating current to the drive motor. It needs either single phase or three-phase AC input. Some power supplies may be external.

Mechanical Package

The disk drive assemblies are housed in the mechanical package, a strong framework often covered with panels.

Summary

Although each disk storage device model is somewhat unique, all disk storage devices have some types of assemblies in common.

Block 9

Magnetic Recording

Recording Fundamentals

This activity introduces some of the disk magnetic recording fundamentals.

Disk system trends and considerations are toward three areas:

- Price
- Throughput
- Capacity

Price is the cost per megabyte of data. Throughput and capacity play important roles in determining the effective cost.

Throughput is the amount of data that can be processed in a given amount of time. The data access time and data flow between the storage device and CPU are the determining factors.

Capacity

Capacity determines much of both price and throughput.

Capacity is the amount of data recorded, and is determined by the number of bits per inch as well as the number of tracks per inch. (See figure 9-1.)

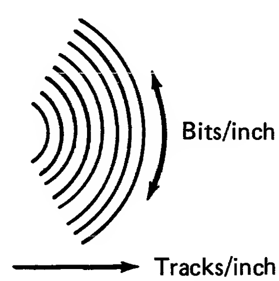

Figure 9-1. TPI/BPI

The capacity of disk systems has been increased by decreasing the following:

- Flying height
- Coating thickness
- R/W core size

Flying Height

As you may recall from magnetic theory, lines of force radiate out from the magnet. Because the lines of force cover a wide area on the recording surface (see figure 9-2), the first thing to do is bring the magnet closer to the surface.

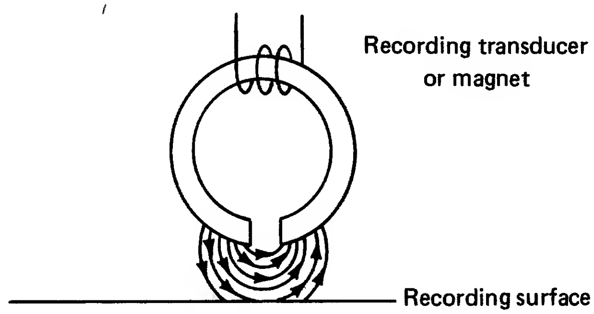

Figure 9-2. Lines of force magnetizing large area

As can be seen, the lines of force will magnetize a smaller area. Thus, more bits can be recorded in a given distance. This also affects the width of the bit, allowing the availability of more tracks per inch (see figure 9-3). However, this is not the primary consideration since it affects fringing; core width affects TPI and is the major consideration

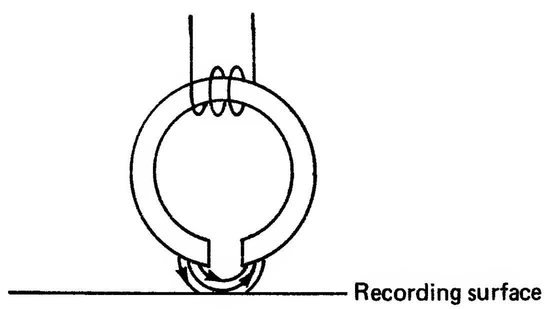

Figure 9-3. Lines of force magnetizing small area

Coating Thickness

When you consider recording fundamentals, you must also consider the coating on disk platters. The early platters were coated with iron oxide approximately 150 microinches thick. When a data bit was written on the media, the residual magnetism could be induced to any depth (see figure 9-4)

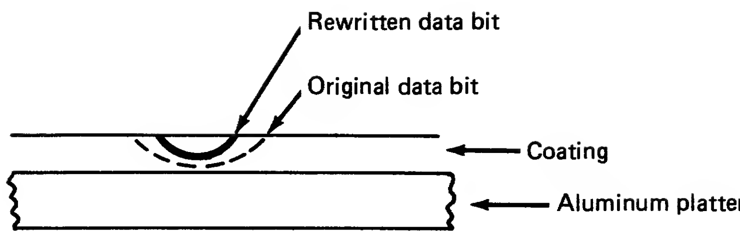

Figure 9-4. Section of disk with data bits

The dotted line represents the original data bit recorded. If information had to be rewritten on the disk, a complete erase might not occur. This could happen due to any of the following factors:

- Different recording device
- Different flying height of write head
- Wobble in the platter
- Weaker write current

As can be seen, incomplete erasure would present data errors because the two signals would interpose or combine. To overcome this, the heads were designed to perform an erase before write; this operation is called a pre-erase. The heads were shaped as shown in figure 9-5.

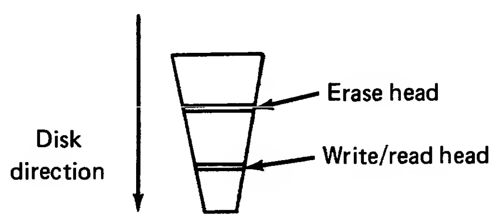

Figure 9-5. Head shape

This method completely erased the old data before the new data was written. During a read operation, erase current was not applied. Pre-erase was an acceptable operation until the track capacity needed to be increased. The erase head was wider than the write head, thus track width was wide. (See figure 9-6.)

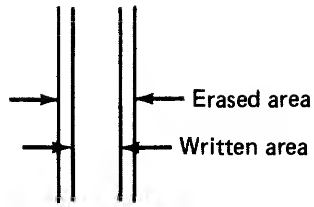

Figure 9-6. Track width

In order to increase track density, it was necessary to eliminate the pre-erase head. To do this and guarantee complete data erase, the disk coating had to be thinner (today's disk coating range is 20 microinches to 50 microinches). Now, the coating could be completely saturated when write current was applied. (See figure 9-7.)

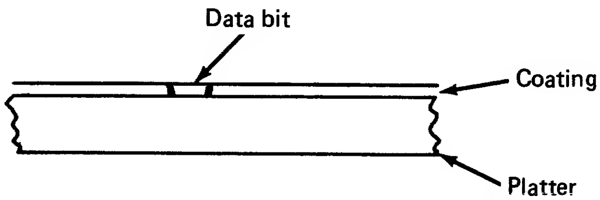

Figure 9-7. Thinner coating

R/W Core Size

Since the recording media was thinner and the heads were flying closer to the media, it became necessary to make the R/W core smaller. The core was made thinner so the track width could not overlap, which causes cross talk, the interference of signals from adjacent tracks. At the same time, because the saturation level of the media was less, the gap in the head was also made smaller. This allowed several things to happen:

- Less current was required for saturation.
- Less data width was required, which resulted in more tracks per inch (TPI).
- Less data length was required, which resulted in more bits per inch (BPI).

Conclusion

The decrease in flying height, coating thickness, and R/W core size enabled increases in the capacity. Early disk drive capacities were about 1000 to 3000 bits per inch with 200 tracks per inch. Today's drives average 4000 to 8000 bits per inch and 400 to 600 tracks per inch.

The smoother disk finish and smaller heads allowed faster rotational speed of the disk; from about 1500 rpm to over 3600 rpm today. Thus, access and data transfer times decreased and the cost per bit decreased.

Problems

The increases in capacity and access times were not accomplished without problems. These included mechanical, environmental, and electromechanical conditions.

Mechanical Conditions

Earlier, you learned how the distance between the head and media was reduced, thus allowing more BPI of data. You might think that the ideal situation would be to have these two in contact; however, remember that the platter is aluminum and rotational speed is in excess of 100 miles per hour. Contact, even for a split second, would generate a high degree of friction heat. This contact, also know as "head crash," is catastrophic. Head crash has been known to cause the following damage:

- Remove coating media from platter surface
- Burn media so magnetic properties are lost
- Demagnetize surrounding areas due to high heat generation
- Destroy heads

The rotational speed of the platter may create some undesirable aerodynamic effects on the disk itself. The disk may ripple, bow up, bow down, or become S shaped. (See figure 9-8.)

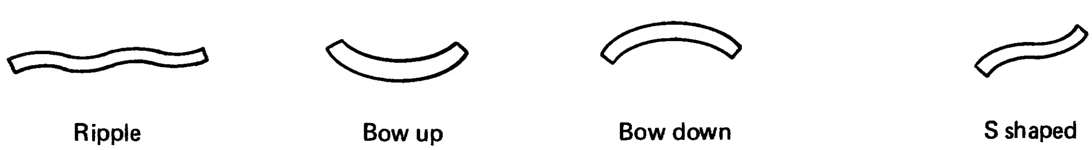

Figure 9-8. Aerodynamic effects on disk

The areodynamic effects may be any one, or a combination, of the shapes. These effects may be caused by improper assembly or foreign matter on the platter mounting surface, unequal air pressure, or defective platter material.

Environmental Conditions

During manufacture of the disk platter, problems may occur if the environmental conditions are not controlled. Some of the problems that could occur are:

- Pits in the disk surface—these may be covered when coated, creating an air pocket that could break the coating
- Impure magnetic coating
- Foreign matter trapped in the coating
- Dust particle scratches on the surface

There are other environmental factors that can influence disk performance. The flying height of the head has an approximate ratio of 1:1 with coating thickness. Thus, flying height = 20 to 50 microinches. A smoke particle is about 250 microinches high. A fingerprint is about 1000 microinches high. A dust particle is about 2500 microinches high. A human hair is about 3000 microinches high. (See figure 9-9.)

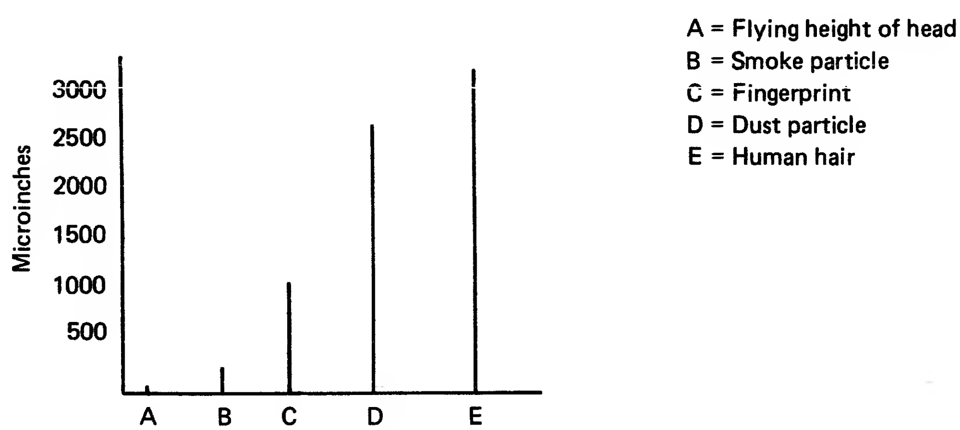

Figure 9-9. Height comparisons

As you can see, the following precautions must be taken when working with disk systems:

- Do not handle or touch disk surfaces.
- Do not smoke around open disk packs.

- Do not blow on the disk surface to remove dust—the moisture expelled by blowing leaves a deposit which is about 500-1200 microinches high when dry.
- Do not attempt to clean disk surfaces without proper equipment.
- Do not attempt to clean disk heads.

As can be seen from figure 9-9, a smoke particle would look like a high mountain to the flying head. Most dust particles are removed from the surfaces by either the forced air system or the head pad before the R/W core hits it. Because the head is floating, the smoke particles are burned off when the head pad crashes into it. Problems do arise with this as previously discussed. Fingerprints and moisture are the most damaging. When dry, both leave a salt and oil residue covering a large surface area. The flying head cannot respond to these; thus, the air bearing is destroyed and the head crashes which destroys the disk surface.

The least amount of damage possible, due to improper care and handling of disk packs, is loss of data. The most damage possible is the destruction of the disk and/or head. Extreme care and caution must be exercised when you are handling or working with disk packs and read/write heads.

Electromechanical Conditions

The last problem area to be aware of in recording fundamentals is electromechanical. This type of problem combines both the electrical and mechanical properties of a high-capacity system. Since no electromechanical device can be perfect, there is an effect that degrades read accuracy by distorting the signal. This effect is called peak shift.

Peak Shift

In an ideal world, the flux reversal would be instantaneous, as shown in the ideal recording portion of figure 9-10. Current would immediately switch from one polarity to the other. As a result, the distance required to complete the magnetic flux reversal on the disk would be so narrow as to be insignificant; the readback pulse would then also be extremely narrow. To carry the principle one step further, the heads would be an extremely small distance from the disk surface. Therefore, the head gap itself could be made very small for two reasons:

1. The magnetic field strength increases as the head moves closer to the disk.
2. The head gap must be wide enough to intersect sufficient lines of force from the magnetic flux field to generate a signal. The weaker the signal, the wider the gap must be. With the substantial flux amplitude gained by having the head very close to the disk surface, a very small head gap can generate a reliable readback voltage.

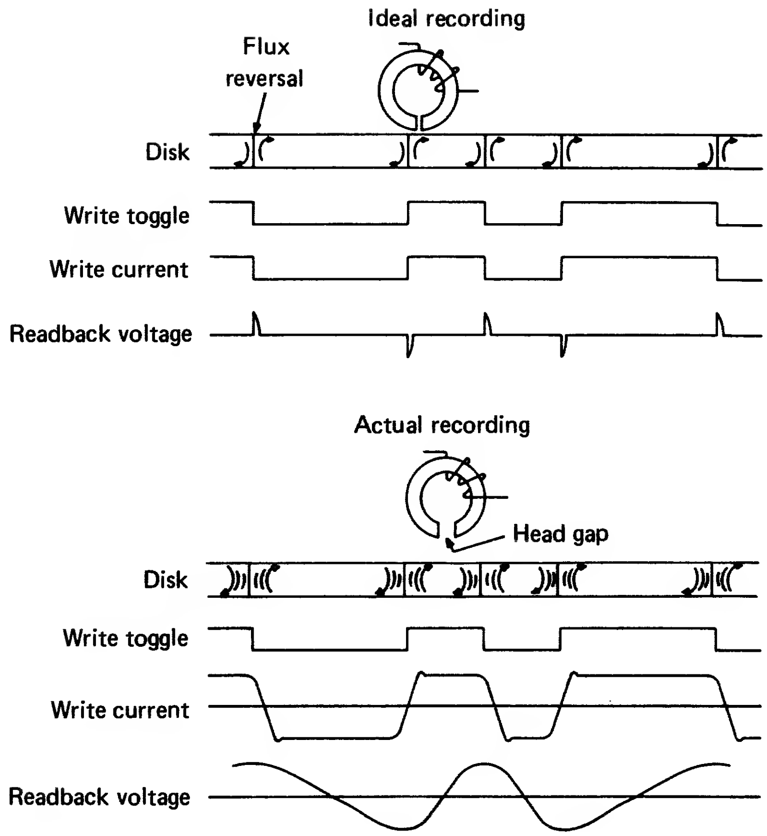

Figure 9-10. Write irregularity

But in the real world, it takes time for the current to reverse; the flux change is not instantaneous. Furthermore, heads must fly a finite distance from the disk. The greater the distance between the head and the oxide, the wider the head gap must be. The resulting readback voltage is more or less sinusoidal with peaks less easily defined in time or amplitude.

With modern high-frequency recording techniques, adjacent clock/data pulses are close enough to intersect with each other. This is shown in figure 9-11. Peak shift is the result of the interaction of the pulses. Because two pulses tend to have a portion of their individual signals superimpose themselves on each other, the actual readback voltage is the algebraic summation of the pulses.

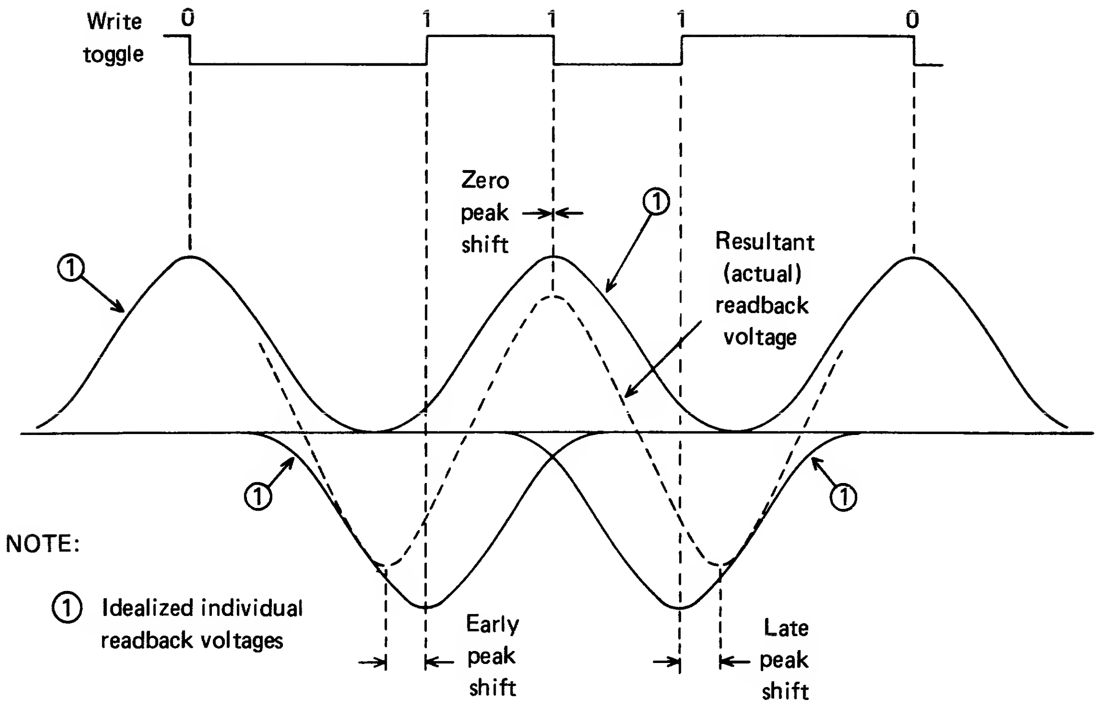

Figure 9-11. Peak shift

When all 1s or all 0s are being recorded, the data frequency is constant: pulses are spaced apart by one cell. As a result, the pulse spacing causes the overlap errors to be equal and opposite. The negative-going and positive-going errors cancel each other. This is the "zero peak shift" condition of the ". . .111. . ." pattern in figure 9-11.

Peak shift occurs when there is a change in frequency. A 011 pattern represents a frequency increase since there is a delay of about 1.5 cells between the 01 and only 1.0 cell between the 11. As a result, the squeezing of the cells causes the mathematical average (the actual readback voltage) to shift the apparent peak to the left. This is early peak shift.

On the other hand, a 10 pattern represents a frequency decrease since a pulse is not written at all in the second cell. In addition, a 001 pattern is also a frequency decrease since there is a 1.0 cell interval between the first two bits and 1.5 cell interval between the last two bits.

These examples examined only two or three bits without regard to the preceding or subsequent data pattern. The actual bit combinations can be considerably more complex. Any data pattern will have considerable overlapping of the data pattern frequency changes. Consider the overlap of these eight bits:

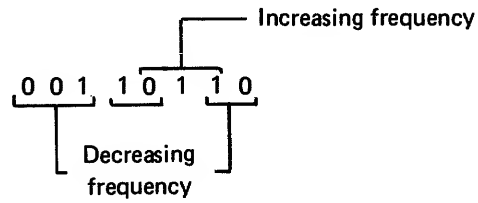

Any of these peak shift conditions, if severe enough, can cause errors during subsequent read operations.

Digital Encoding Techniques

There are many codes and formats used in magnetic recording. The codes are selected for their ability to meet the objective of low error rate at high-density recording. To meet this objective, codes must have the following standards:

- Little or no DC component
- Narrow bandwidth
- Low transition rate-to-bit-rate ratio

A DC component within the signal causes a shift in the base line of the read circuitry; this limits the high-density efficiency.

Wide bandwidths also cause peak shifts due to media and electronic noise; this limits the high-density efficiency.

The newer, self-clocking codes require enough transitions, within a random data sequence, to allow the clock to be obtained from these transitions.

This activity discusses various recording codes and formats that have been or are being used. Special attention should be paid to the section on modified frequency modulation (MFM) as this is the technique presently being used by most disk manufacturers. The advantages and disadvantages of all techniques are covered. Familiarize yourself with the following recording techniques:

- Return to zero (RZ)
- Double pulse return to zero (DPRZ)
- Return to zero with prebias (RZ with PB)
- Return to negative (RTN)
- Nonreturn to zero (NRZ)
- Pulse envelope recording (PER)
- Phase modulation (PM)
- Frequency modulation (FM)
- Modified frequency modulation (MFM)

RZ—Return to Zero (Discrete Pulse)

See figure 9-12.

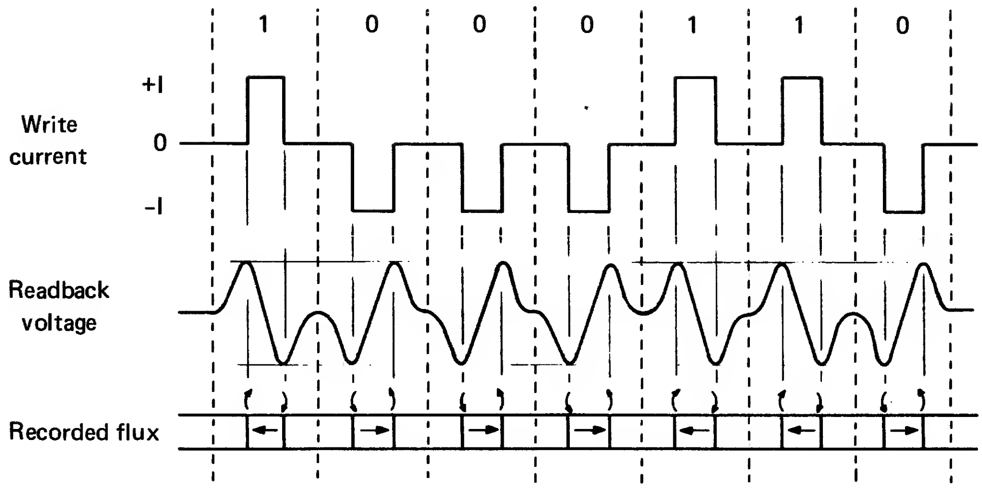

Figure 9-12. Return to zero

Logical Rules for Recording

A 1_2 is recorded by a positive current pulse, and a 0_2 is recorded by a negative current pulse.

Record/Playback Logic

Record/playback logic for RZ recording is shown in figure 9-13.

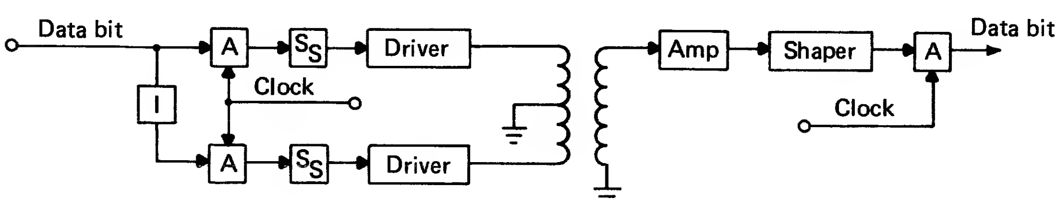

Figure 9-13. Record/playback logic

Readback Recognition of Bits

A 1_2 is recognized by a positive followed by a negative voltage pulse within a cell; a 0_2 is recognized by a negative followed by a positive voltage pulse within a cell. A 180 degree phase difference is apparent between 1 and 0 pulses.

Advantages

The following list itemizes the advantages of RZ recording.

- Ability to selectively alter (write) or read by bit
- Less susceptible to noise produced by recording surface defects, due to neutral magnetization between bits
- Does not require a lot of circuitry; about same as NRZI
- Low average head writing current is realized due to a low-duty cycle
- Self-clocking and self-checking since two pulses always appear in every cell, (lack of signal is obviously an error)
- A 180 degree phase difference, rather than amplitude, distinguishes 1 and 0 pulses

Disadvantages

The following list enumerates the disadvantages of RZ recording.

- Pulse packing density limit is less than that of NRZ and NRZI since two flux changes (instead of one as for NRZI) are contained within a cell.
- External means must be provided to pre-erase old information before recording new if self-clocking system is employed (requires neutralization of medium between bits); pre-erasing is generally not necessary if a clocking track is provided which identifies and fixes each cell position.
- About one-half read amplitude of NRZ and NRZI because recording is from neutral to saturation rather than from saturation to opposite saturation.

DPRZ—Double Pulse Return to Zero

See figure 9-14.

Logical Rules for Recording

A 1_2 is recorded by a positive rectangular current pulse followed by a negative pulse one-half of a celltime later. A 0_2 is recorded by a negative current pulse followed by a positive pulse one-half of a celltime later. If pulse width is extended to one-half of a celltime, DPRZ becomes identical to PM recording.

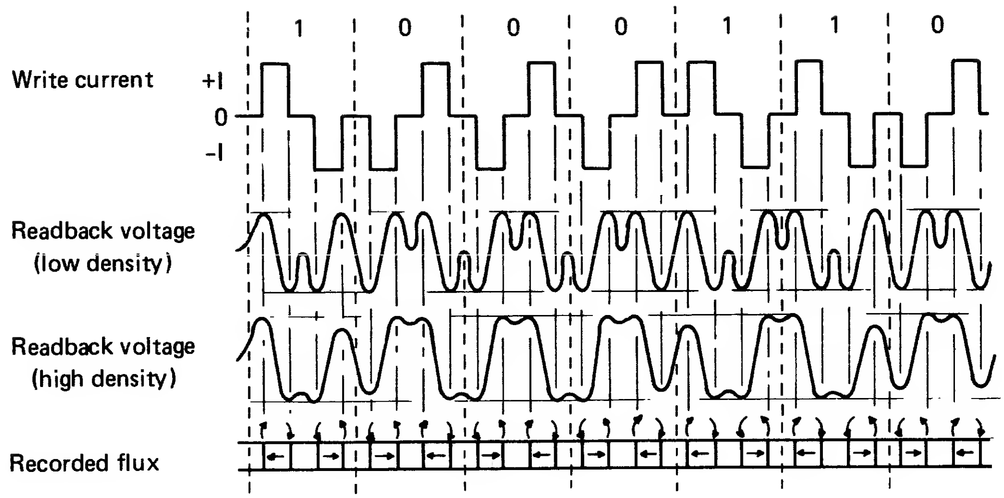

Figure 9-14. Double pulse return to zero

Record/Playback Logic

Figure 9-15 shows record/playback logic for DPRZ recording.

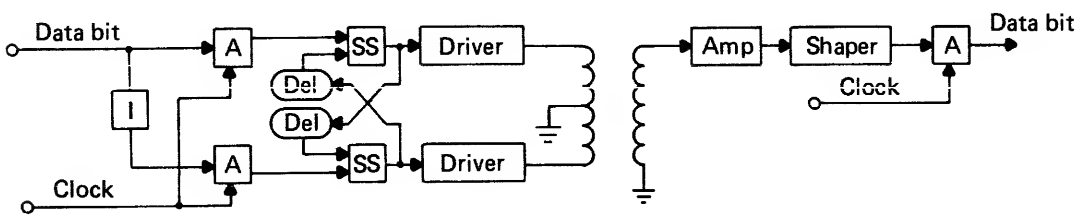

Figure 9-15. Record/playback logic (DPRZ)

Readback Recognition of Bits

A 1_2 is characterized by a negative-going voltage crossover followed by a positive-going voltage crossover one-half of a celltime later. A 0_2 is characterized by a positive-going voltage crossover followed by a negative-going voltage crossover one-half of a celltime later.

Advantages

The following list enumerates the advantages of DPRZ recording.

- It has been reported that an increase in bit resolution is obtained over the RZ method. With the RZ method, a string of 0s or 1s closely packed together results in enough flux crowding to merge playback pulses resulting in a very low resolution. The DPRZ method appears to compensate for this effect due to the appearance of oppositely polarized flux excursions within a group of 0s or 1s.
- Average current flowing through head is reduced due to low duty cycle, which may be important if any dissipation or heating effects become a problem.

Disadvantages

The following list describes the disadvantages of DPRZ recording.

- Resolution of packing density limit is less than NRZI but greater than RZ.
- The medium must be AC erased prior to recording if a self-clocking system is employed (requires neutralization between bits). Pre-erasing is generally not necessary if a clocking track is provided which identifies each cell position.

RZ With PB—Return to Zero with Prebias (Prebiased Discrete Pulse Recording)

See figure 9-16.

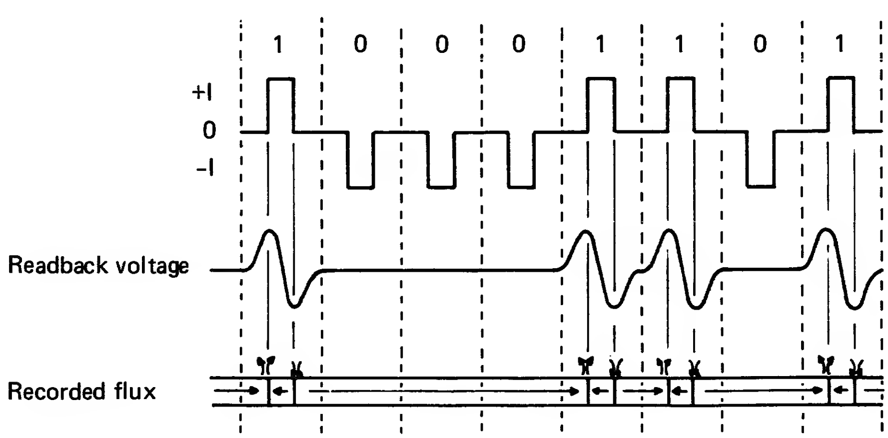

Figure 9-16. Return to zero with prebias

Logical Rules for Recording

A 1_2 is recorded by a positive rectangular current pulse at the center of a cell. A 0_2 is recorded by a negative rectangular current pulse (or by the absence of a pulse at the center of a cell). Recording must be on a medium that is prebiased to a negative saturation level.

Record/Playback Logic

Figure 9-17 shows record/playback logic for RZ with PB recording.

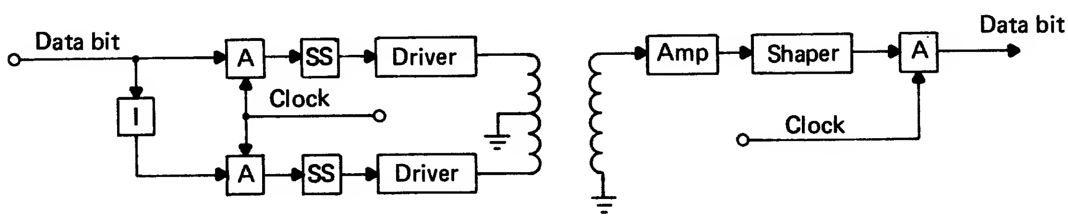

Figure 9-17. Record/playback logic (RZ with PB)

Readback Recognition of Bits

Because the recording is on a medium prebiased to a negative saturation, the negative pulses used for recording 0s will not change the prebias state. However, positive current pulses, used for recording 1s, will drive the medium to the positive state. Therefore, during the process of reading back the information, 0s will be characterized by an absence of a pulse, and 1s will be characterized by a positive followed by a negative voltage pulse within a cell.

Advantages

The following list enumerates the advantages of RZ with PB recording.

- The process of prebiasing or erasing is a little simpler than in the conventional RZ method. Instead of erasing with high-frequency erase current, as with RZ, erasing can be done with either an electromagnet, permanent magnet or just another write head driven with a constant DC bias current.
- The READBACK signal amplitude is about twice that of RZ.
- The READ signal is always the same shape and contains the same basic frequency. This permits noise elimination by use of a narrow bandpass amplifier.

Disadvantages

The following list describes the disadvantages of RZ with PB recording.

- Pulse packing density limit is less than NRZI since two flux changes are contained within a cell.
- Determining the difference between a drop-out and a zero is impossible unless a parity bit or other checks are used.
- Readback noise, due to imperfections in the medium, is at a maximum since the medium is saturated in the intervals between pulses. This effect, however, can be minimized by the use of high-quality recording media.

RTN—Return to Negative (Return to Saturation; Biased Discrete Pulse Recording)

See figure 9-18.

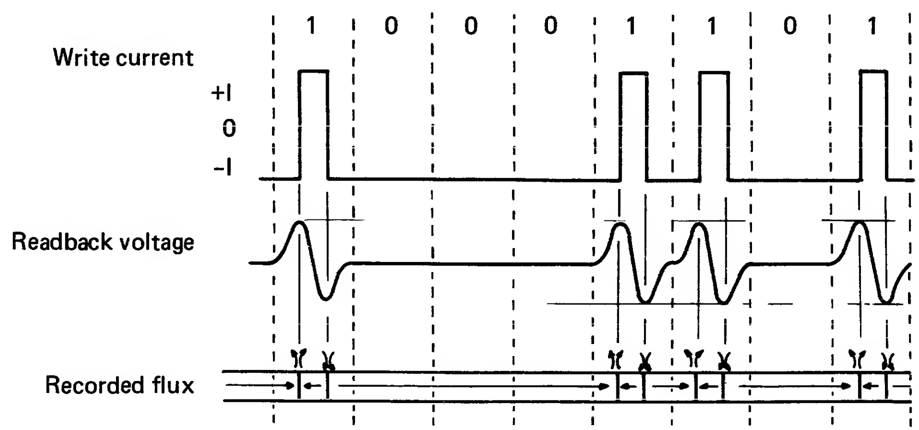

Figure 9-18. Return to negative

Logical Rules for Recording

A 1_2 is recorded by a positive rectangular current pulse at the center of a cell period. A 0_2 is recorded by the absence of a pulse. Write current is always returned to a negative level following the writing of a 1.

Record/Playback Logic

Figure 9-19 shows record/playback logic for RTN recording.

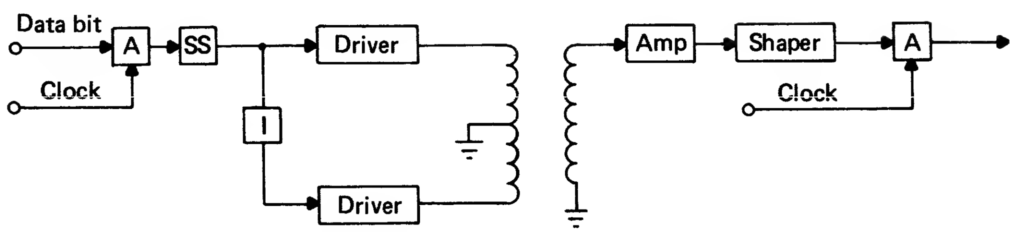

Figure 9-19. Record/playback logic (RTN)

Readback Recognition of Bits

Upon read back, 1s will be represented by a positive followed by a negative voltage pulse within a cell period and 0s will be characterized by the absence of a pulse. This method is very similar to RZ with PB since the same recorded flux and readback voltage waveforms are obtained.

Advantages

The following list enumerates the advantages of RTN recording.

- No prebias (erase) operation is necessary since recording is done at saturation, excluding other factors.
- The READBACK signal amplitude is about twice that of RZ.
- The READ signal is always the same shape and contains the same basic frequency. This allows noise elimination by use of narrow bandpass amplifier.

Disadvantages

The following list describes the disadvantages of RTN recording.

- Pulse packing density limit is less than NRZI since two flux changes are contained within a cell.
- Differentiation between a dropout and a zero is impossible unless a parity bit or other checks are used.
- Readback noise due to imperfections in the medium is at a maximum since the medium is saturated in the intervals between pulses. This effect, however, can be minimized by the use of high-quality recording media.

NRZ—Nonreturn to Zero

See figure 9-20.

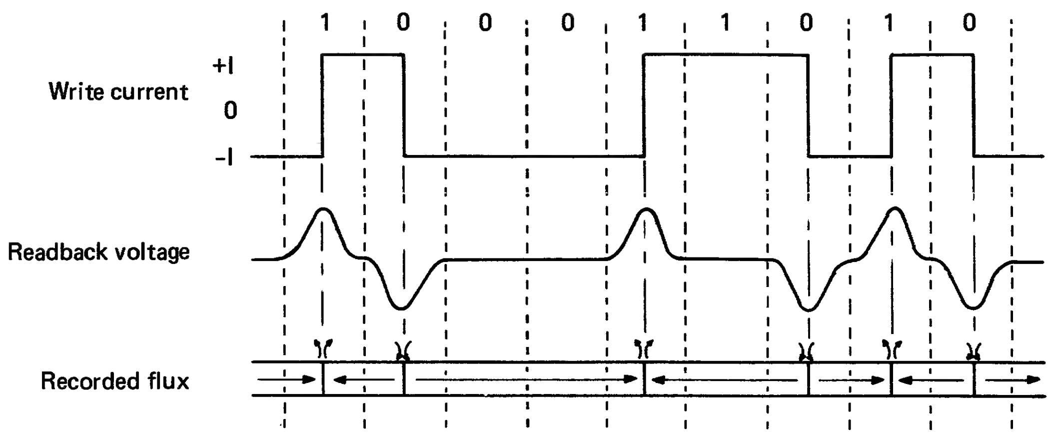

Figure 9-20. Nonreturn to zero

Logical Rules for Recording

A 1_2 is recorded by a positive current throughout the entire bit period. A 0_2 is recorded by a negative current throughout the entire bit period.

Record/Playback Logic

Figure 9-21 shows the logic for NRZ record/playback.

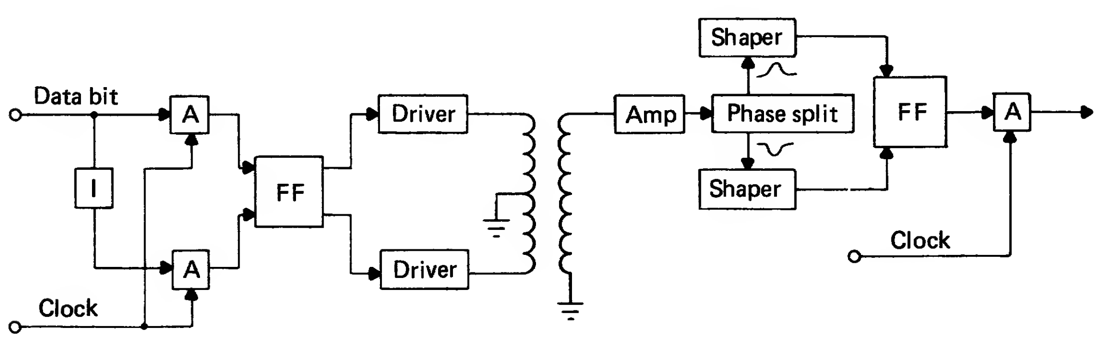

Figure 9-21. Record/playback logic (NRZ)

Readback Recognition of Bits

A 0 cell followed by a 1 cell is characterized by a positive readback pulse. A 1 cell followed by a 0 cell is characterized by a negative readback pulse. A string of 1s is characterized by an initial positive pulse followed by the absence of a pulse for every 1 cell in that string. Similarly, a string of 0s is characterized by an initial negative pulse followed by an absence of a pulse for every 0 cell in that string.

Advantages

The following list enumerates the advantages of NRZ recording.

- The primary advantage of NRZ is in realizing a high bit packing capability from any medium since a maximum of only one flux change per bit is required.
- Since recording is done at saturation, the medium need not be erased prior to recording, excluding other factors.
- Readback signal amplitude is about twice that of RZ.

Disadvantages

The following list describes the disadvantages of NRZ recording.

- The reading process is a little more complex since it is necessary for the reading circuit to remember the polarity of the first bit of a string of identical bits in order to identify these bits.
- Some form of read clocking is necessary to determine when to sample for the presence or absence of a signal. An initial starting condition must be defined because the first string of bits in a record may be recorded at the same flux direction as the area before the record resulting in no initial pulse for that string.
- Differentiation between a dropout and an intentional lack of a READ signal is impossible unless a parity check or other checks are used. A dropout could result in a number of consecutive errors.
- Readback noise, due to imprefections in medium, is at a maximum since the medium is saturated between pulses. This can be minimized by the use of high quality recording media.
- NRZ requires higher power handling capability in write head and write drivers than in the case of RZ, RZ with PB, and DPRZ.
- Read amplifier must have wide frequency range.
- NRZ requires a little more read circuitry than NRZI.

PER—Pulse Envelope Recording

See figure 9-22.

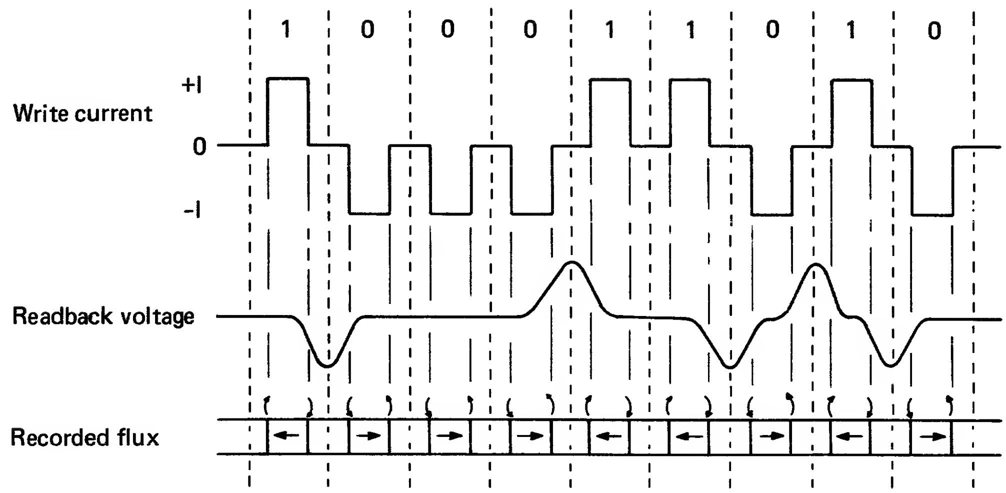

Figure 9-22. Pulse envelope recording

Logical Rules for Recording

A 1_2 is recorded by a positive current pulse. A 0_2 is recorded by a negative current pulse.

Record/Playback Logic

Figure 9-23 shows record/playback logic for PER recording.

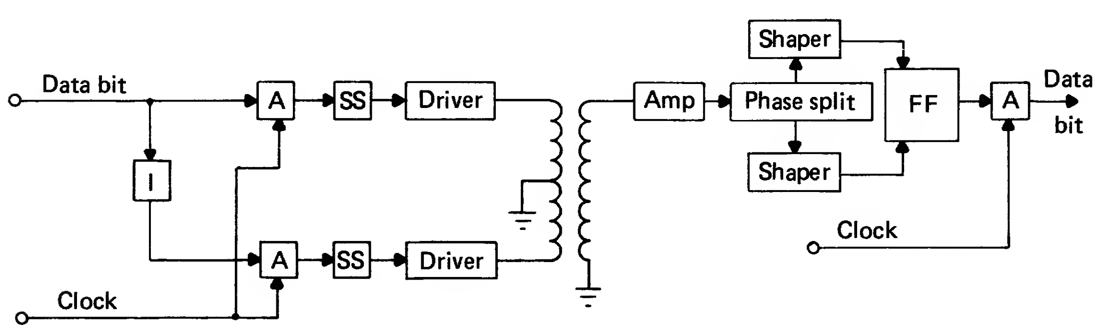

Figure 9-23. Record/playback logic for PER

Readback Recognition of Bits

In pulse envelope recording, discrete pulses are recorded at such a high density that head resolution in both reading and writing causes only 1 to 0 and 0 to 1 changes to read back. Thus the readback is in the form of NRZ.

Advantage

An advantage of PER recording is that it uses a lower duty cycle than NRZ.

Disadvantage

A disadvantage of PER recording is the double frequency requirements in the write circuits.

NRZI—Nonreturn to Zero Indiscrete (Pouliart's Variation)

See figure 9-24.

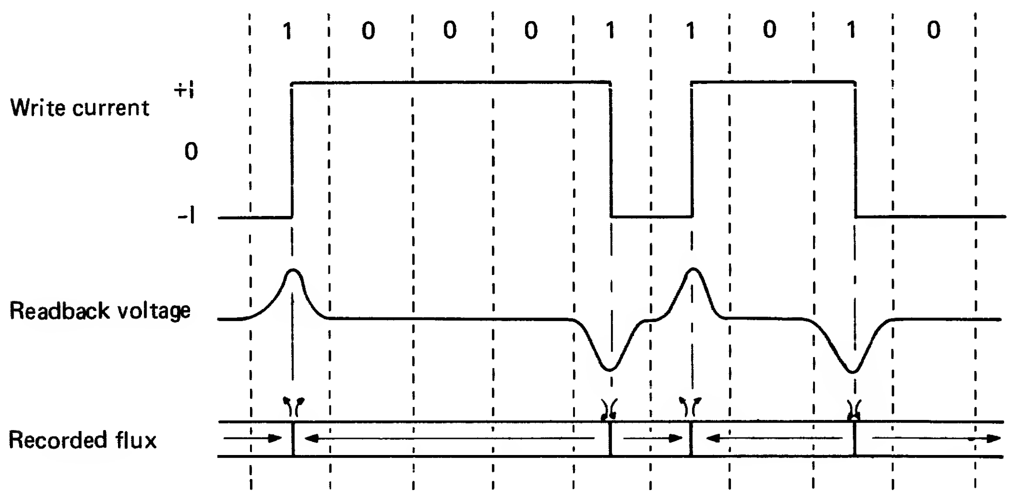

Figure 9-24. Nonreturn to zero indiscrete

Logical Rules for Recording

A 1_2 is recorded by a reversal of current direction. A 0_2 is recorded by no reversal of current.

Record/Playback Logic

Figure 9-25 shows record/playback logic for NRZI recording.

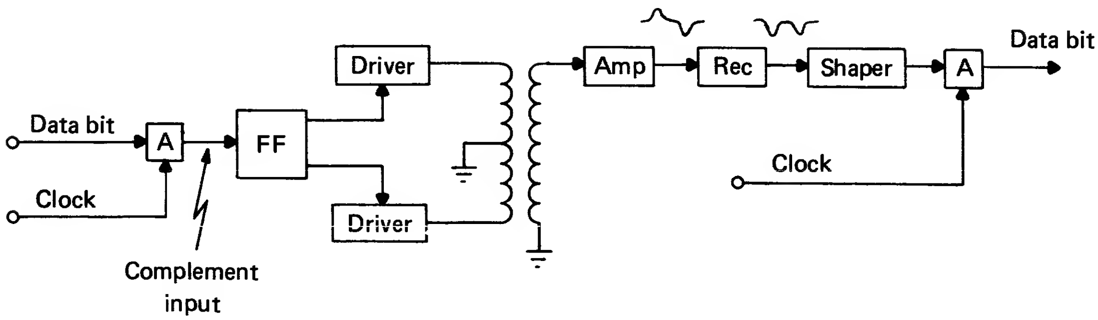

Figure 9-25. Record/playback logic for NRZI

Readback Recognition of Bits

Since a flux reversal occurs for every 1 written, only 1s will appear as pulses, regardless of polarity. The absence of a signal identifies 0s.

Advantages

The following list enumerates the advantages of NRZI recording.

- High bit packing capabilities can be realized since a maximum of only one flux change is required per bit.
- Since recording is done at saturation, the medium need not be erased prior to recording, excluding other factors.
- NRZI uses simpler readback logic circuitry than NRZ.
- Signal amplitude is about twice that of RZ methods.
- A dropout during playback will result in only one error rather than a string of errors which is possible with NRZ.

Disadvantages

The following list defines the disadvantages of NRZI recording.

- It is not self-clocking because information is recorded by the absence of a pulse during readback. This requires use of a clock track or schemes using other tracks or oscillators.
- It is impossible to differentiate between a dropout or intentional absence of a pulse without use of more elaborate checking schemes, e.g., parity checking.
- Read-back noises, due to imperfections in medium, are maximum since the medium is saturated between bits. This effect can be minimized by use of high quality recording media.
- NRZI requires greater power handling capability in write head and drivers than RZ methods.
- Read amplifier must handle a wide range of frequencies.
- Most computer magnetic tape applications can use NRZI recording.
- Some disk file applications, e.g., IBM® 1311 Disk Pack, can use NRZI recording.

PM—Phase Modulation (MNRZ—Modified Nonreturn to Zero; William's Phase Modulation; Manchester Recording; Double Transition; Ferranti Recording)

See figure 9-26.

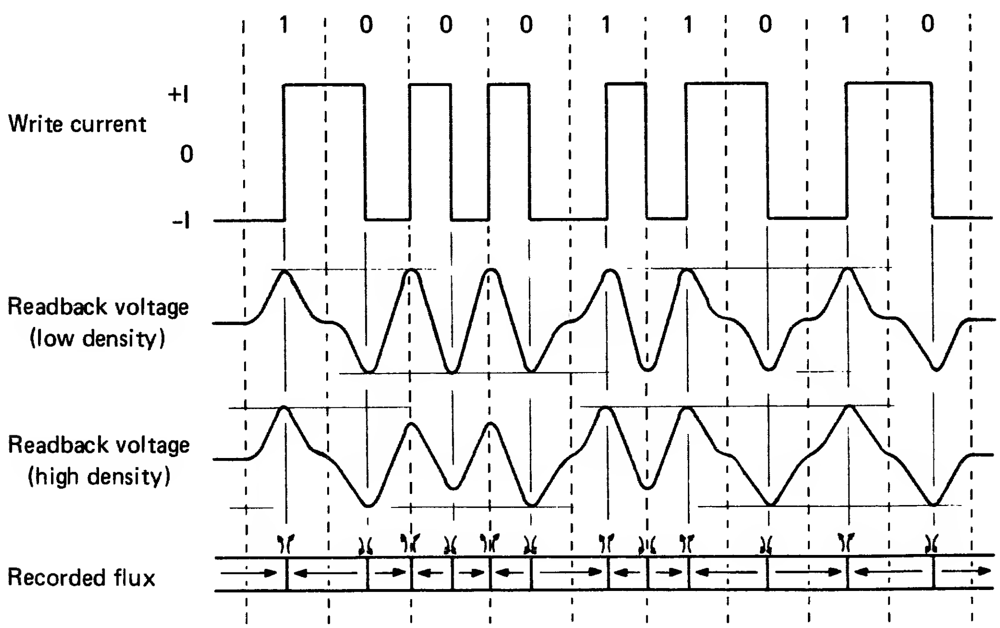

Figure 9-26. Phase modulation

Record/Playback Logic

Figure 9-27 shows record/playback logic for PM.

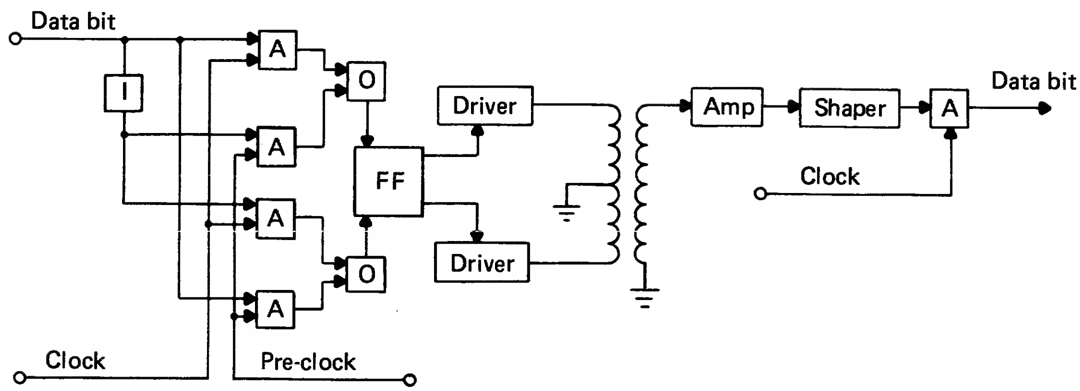

Figure 9-27. Record/playback logic (PM)

Logical Rules for Recording

A 1_2 is recorded by a positive current change. A 0_2 is recorded by a negative current change. This necessitates a negative current change between consecutive 1 cells and a positive current change between consecutive 0 cells.

Readback Recognition of Bits

A 1_2 is identified by a positive pulse at the center of a cell. A 0_2 is identified by a negative pulse at the center of a cell. Other characteristics of the READBACK signal are apparent upon examination of the waveform. These characteristics are as follows:

- A positive pulse at the junction of two cells indicates that both cells contain 0s. Conversely, a negative pulse at the junction of two cells indicates that both cells contain 1s.
- The absence of a pulse at the junction of two cells indicates that the cells contain opposite binary bit values.

Thus, the presence or absence of a pulse at the junction of two cells not only can predict the succeeding binary bit but also serves as a checking feature on a bit-by-bit basis.

Advantages

The following list describes the advantages of phase modulation recording.

- Clocking and checking may be simplified since a pulse always appears at the center of a cell.
- The binary bit can be identified by polarity rather than amplitude.
- The frequency response requirement of the overall record/playback system is relaxed somewhat because of the maximum flux to transition period of two to one (i.e., limited to one octave).
- The medium need not be erased prior to writing, excluding other factors.
- Signal amplitude is about twice that of RZ methods.
- Theoretical effective signal amplitude is twice that of NRZI. This is due to the fact that both ones and zeros are recorded, and the difference between a one and a zero is twice as large at strobe time during readback. This implies a signal-to-noise ratio of 2 to 1 in favor of PM over NRZI as long as the density is one-half NRZI.
- Higher bit densities than NRZI (1.5 times maximum NRZI) are possible at the expense of signal-to-noise ratio.

Disadvantages

The following list describes the disadvantages of phase modulation recording.

- Write circuitry is about twice as complex as NRZI.
- Large signal amplitude variations are encountered when reading at densities near or beyond maximum possible NRZI density.
- Readback noise due to imperfections in the medium are at a maximum since the medium is always saturated. This effect can be minimized by use of high-quality recording media.
- PM recording requires greater power handling capability in write drivers and heads than RZ methods.

FM—Frequency Modulation (Frequency Shift)

See figure 9-28.

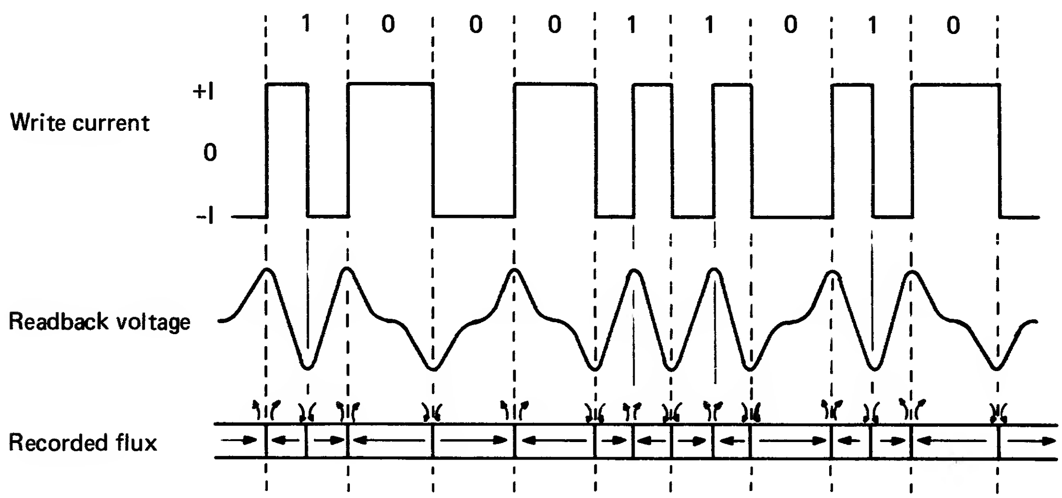

Figure 9-28. Frequency modulation

Logical Rules for Recording

A 1_2 is recorded by a reversal of current at the center of a cell. A 0_2 is recorded by a constant polarity of current for the entire cell period. There is always a current reversal at the beginning and end of each cell.

Record/Playback Logic

Record/playback logic for FM recording is shown in figure 9-29.

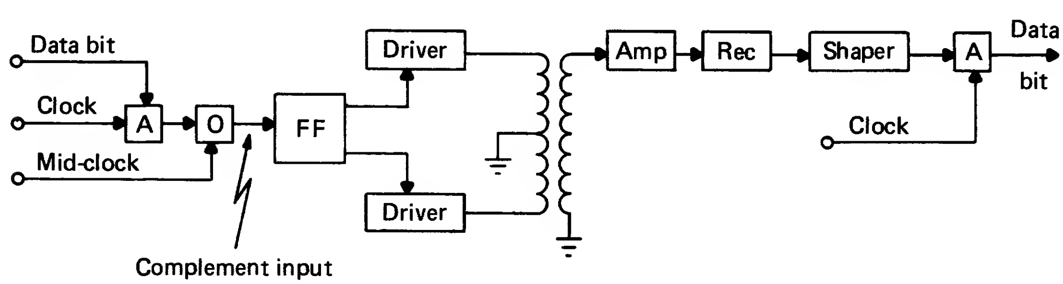

Figure 9-29. Record/playback logic (FM)

Readback Recognition of Bits

A 1_2 is identified by two pulses occurring during a cell period. A 0_2 is identified by only one pulse occurring during a cell period.

A binary bit can also be identified by strobing the pulse polarity at the end of each cell period and comparing it to the polarity at the end of the previous cell. Similar results indicate a stored 1 and dissimilar results indicate a stored 0.

Advantages

The following list enumerates the advantages of FM recording.

- Clocking and checking can be simplified since a pulse (either positive or negative) always appears at each cell junction.
- The frequency response requirement of the overall record/playback system is relaxed somewhat because of the maximum to minimum flux transition period of 2 to 1 (i.e., limited to one octave).
- The medium need not be erased prior to writing since the surface is always saturated by the write head. This excludes other factors which might dictate erasing.
- Signal amplitude is about twice that of RZ methods.
- FM recording does not require as much write circuitry as PM recording.

Disadvantages

The following list describes the disadvantages of FM recording.

- FM requires more complex write and read circuitry than most other methods.
- Large signal amplitude variations are encountered when reading at densities approaching maximum possible NRZI density.
- Pulse polarity has no relation to the value of a bit. Therefore, FM has half the effective signal to noise ratio of PM. This limits the bit density to some value less than PM.
- Readback noise due to imperfections in medium are maximum since the medium is always saturated. This effect can be minimized by use of high-quality recording media.

MFM—Modified Frequency Modulation (Miller, DMM—Delay Modulation Mark)

MFM code is the format generally used for disk recording today. See figure 9-30.

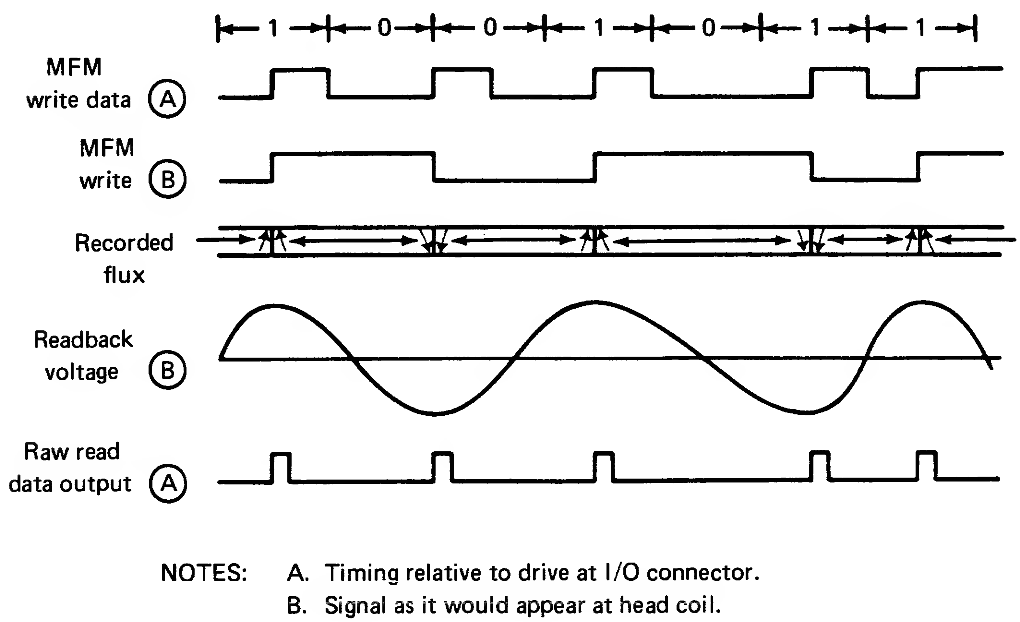

Figure 9-30. MFM recording

Logical Rules for Recording

MFM defines a 1 by writing a pulse at the half-cell time (figure 9-30). A 0 is defined by the absence of a pulse at the half-cell time. A pulse at the beginning of a cell is clock; however, clock is not always written. Clock is suppressed if there will be a 1 in this cell or if there was a 1 in the previous cell.

The rules for MFM recording may be summarized as follows:

1. There is a flux transition for each 1 bit at the time of the 1.
2. There is a flux transition between each pair of 0 bits.
3. There is no flux transition between the bits of a 10 or 01 combination.

Advantages

Fewer flux reversals are needed to represent a given binary number because there are no flux reversals at the cell boundaries. This achieves higher recording densities of data without increasing the number of flux reversals per inch.

Signal-to-noise ratio, amplitude resolution, read chain operation, and operation of the heads are improved by the lower recording frequency achieved because of fewer flux reversals required for a given binary number.

Disadvantages

Pulse polarity has no relation to the value of a bit without defining the cell time along with cell polarity. This requires additional read/write logic and high-quality recording media to be accomplished.

Comparison

Figure 9-31 is a comparison chart of recording techniques.

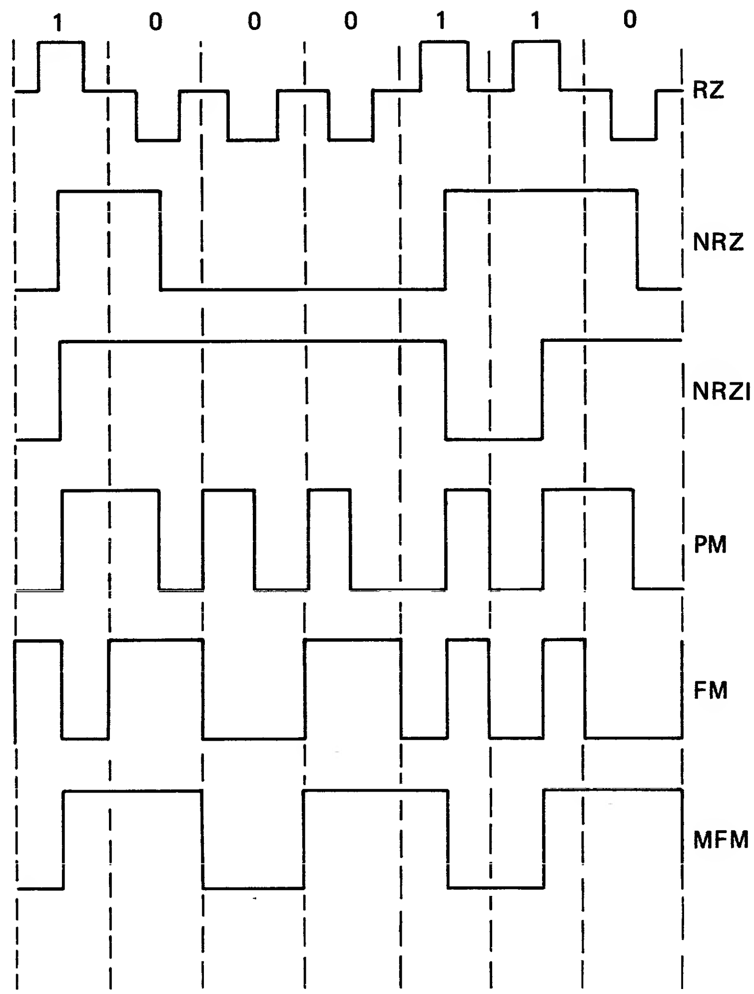

Figure 9-31. Comparison of signals

Signal Processing

This activity introduces some of the concepts necessary to read or write data on a disk. You will look at the layout of a disk pack and disk surface, and you will learn how address identification is assembled on the disk.

A short review of the basic read/write principles will be provided to refresh your memory as to why a write compensation network is necessary. A brief introduction to the write compensation circuit will explain how peak shift is eliminated and the reason for current zoning.

The last portion of this activity introduces the read block diagram and explains how data is read and transferred.

Introduction

An access arm for a disk contains one or more read/write heads. If the access arm has two read/write heads, there is one for the top and one for the underside of a disk (two different disks). To read or write a record, an access arm must be positioned on the disk over the location to be used. The arm (if a moving head) moves in or out to locate itself over the correct spot. If there is not an arm for each disk, the arm must move out from the stack, up or down and in on the correct disk. Having more than one arm reduces the average time required to read or write a record, since access time depends on how far the arm must move from the previous position and how far the disk surface must rotate. A read/write head for each disk surface means that a cylinder of tracks extending vertically through the stack can be accessed (see figure 9-32). Some disk units are available with a read/write head for each track. This arrangement eliminates arm movement altogether.

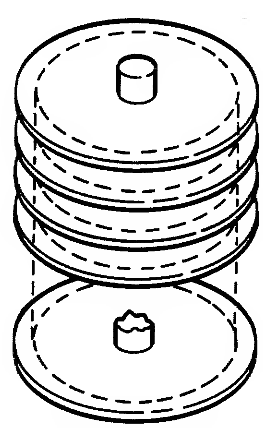

Figure 9-32. Cylinder concept of data recording

Disk Surface and Address Identification

Storage locations on the magnetic disk have an address. For addressing purposes, the disks are identified by disk surface (head), track (cylinder) on the disk surface, and sectors of the track (see figure 9-33). By means of these three identifiers, each storage segment of the disk has a unique address. This entire segment is usually read or written by a READ or WRITE command; alternately, an instruction may read or write an entire track rather than a sector on the track.

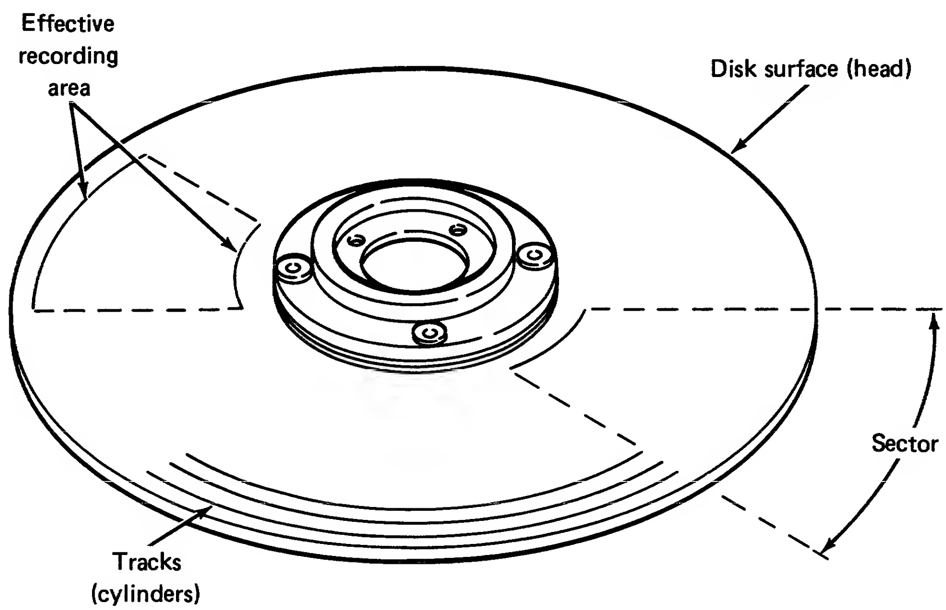

Figure 9-33. Disk address identification

The disk is rotated about its cylindrical axis and the head is held in close proximity to the flat surface of the disk. The data is stored in concentric tracks on the surface. The number of tracks varies with the diameter of the disk. Notice that the tracks are longer near the outer edge of the disk than they are closer to the center. However, the outer tracks do not store more data. All sectors contain the same amount of data regardless of the length of the sector. Programming disk input/output commonly involves the following instructions:

1. A seek instruction to position the access arm at the proper track before reading or writing.
2. A head select instruction to select the desired disk on multiple disk packs.

3. A read instruction to read a record from the disk into an area in main memory, or a write instruction to write a record from main memory onto the disk.
4. Instructions for error detection and control. An example of an error control instruction is an instruction to read the record just written and compare it with the record contained in main memory in order to detect recording errors.

   The disk system functions mainly as a large quantity of medium speed access add-on storage. The disk pack (removable stack) type devices can be used as the system data bank similar to the way in which magnetic tape is used. The fixed disk (permanent stack) is primarily used as temporary storage for often-used programs. Both types of disks are used from time to time for the writing of miscellaneous data; the previously mentioned uses (data bank and temporary storage) overlap on the two types to some degree. Another, less often used, application is as a data transfer vehicle between two computer systems. This is possible if a multiaccess controller (synchronizer) is available for the disk system.

Basic Read/Write Principles

Information is recorded on, and read from, the disk pack by means of read/write heads. Each head contains a read/write coil and, in some instances, an erase coil.

Data is written by passing a current through a read/write coil within the selected head. This generates a flux field across the gap in the head (see figure 9-34). The flux field magnetizes the iron oxide particles bound to the disk surface. Each particle is then the equivalent of a miniature bar magnet with a north pole and a south pole. The writing process orients the poles to permanently store the direction of the flux field as the oxide passes beneath the head. The direction of the flux field is a function of write current polarity while its amplitude depends upon the amount of current: the greater the current, the more oxide particles are affected. This is also known as saturation recording, and is also dependent on optimum timing of the write signal, media thickness, and head flying height.

Information (data) is written by reversing the current through the head. This change in current polarity switches the direction of the flux field across the gap. The flux change defines a data bit.

On those systems equipped with an erase coil, old data is erased by passing a constant current through the erase coil in the read/write head. The constant current, since there are no flux reversals, destroys all old data. The erase gap is wider than the read/write gap and is positioned immediately before the read/write gap. This ensures complete erasure of old data and minimizes crosstalk between adjacent data tracks.

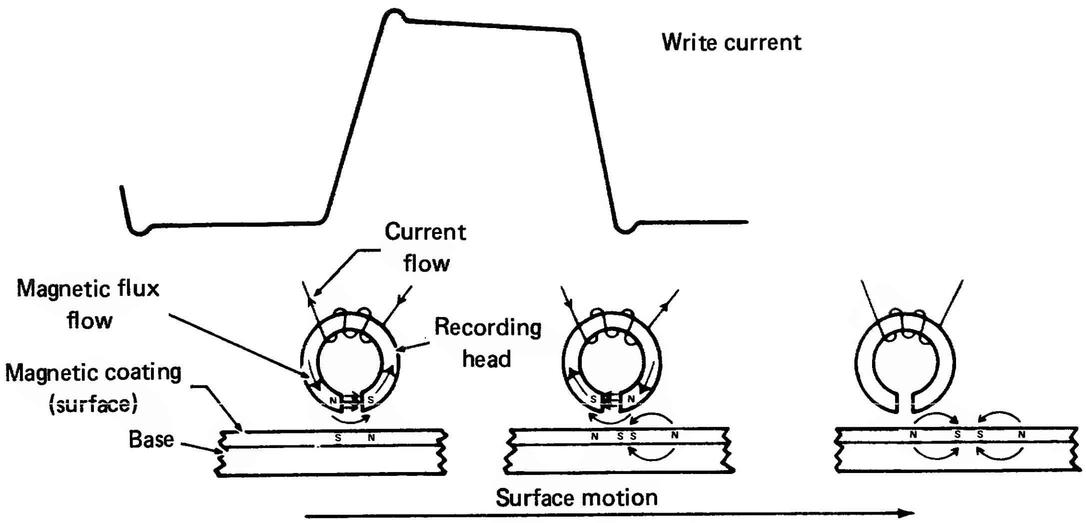

NOTE: Relative head to surface motion, recording (write operation)

Figure 9-34. Writing data

How Data is Retrieved

As the disk passes beneath the read/write head, the stored flux intersects the gap (see figure 9-35). Gap motion through the flux induces a voltage in the head windings. This voltage is analyzed by the read circuit to define the data recorded on the disk.

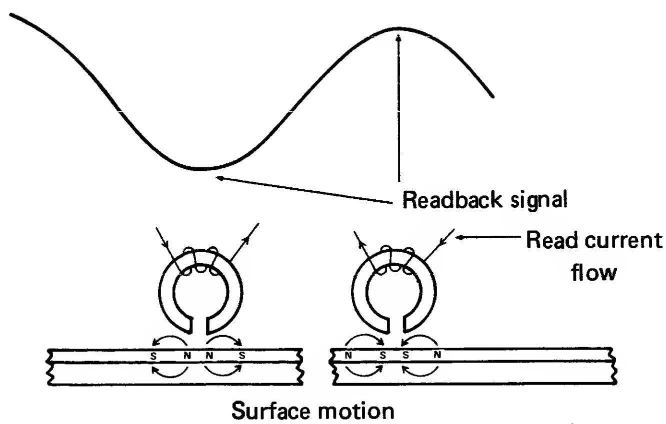

NOTE: Relative head to surface motion, reproducing (read operation)

Figure 9-35. Reading data

Each flux reversal (caused by a current polarity change while writing) generates a readback voltage pulse. Each pulse, in turn, represents a data bit.

The high-density read waveform is somewhat sinusoidal due to a fringing field effect. As you see in figure 9-35, as the media is rotated past the head gap, more and more lines of force are cut. When the maximum number of lines of force are cut (under the gap), the peak of the waveform is reached. It should also be obvious that the closer the head is to the surface, the more lines of force are cut, the higher and sharper the peak. The further the head is from the surface, the less lines of force cut, and the flatter the waveform (READ signal).

Write Compensation

As the capacity of disk systems increases, so does the problem of being able to write and read data without errors. Several examples of problems encountered are:

- Data pattern changes (peak shift). As you may recall, the data frequency (bit pattern recorded versus cell width) is constant whenever all 1s or all 0s are being recorded because all pulses are seperated by one cell. However, as with MFM recording, a 011 pattern represents a frequency increase since there is a delay of about 1.5 cells between the 01, and only one cell between the 11. A 001 pattern represents a frequency decrease since there is only one cell between the 00, and 1.5 cells between the 01.
- Disk rotational speed. The effective linear speed of the disk is smaller near the center than the edges. This, therefore, reduces the air pressure that holds the flying head up. This means that the flying height of the head is closer to the surface near the center. If the write current is held constant over the entire surface, not enough media saturation occurs at the outer edges, and oversaturation occurs near the center.

Write Compensation Circuit

In order to compensate for the problems of peak shift, a pattern decode circuit is installed in the disk drive which decodes the input data. A simplified block diagram of the write compensation circuit is shown in figure 9-36.

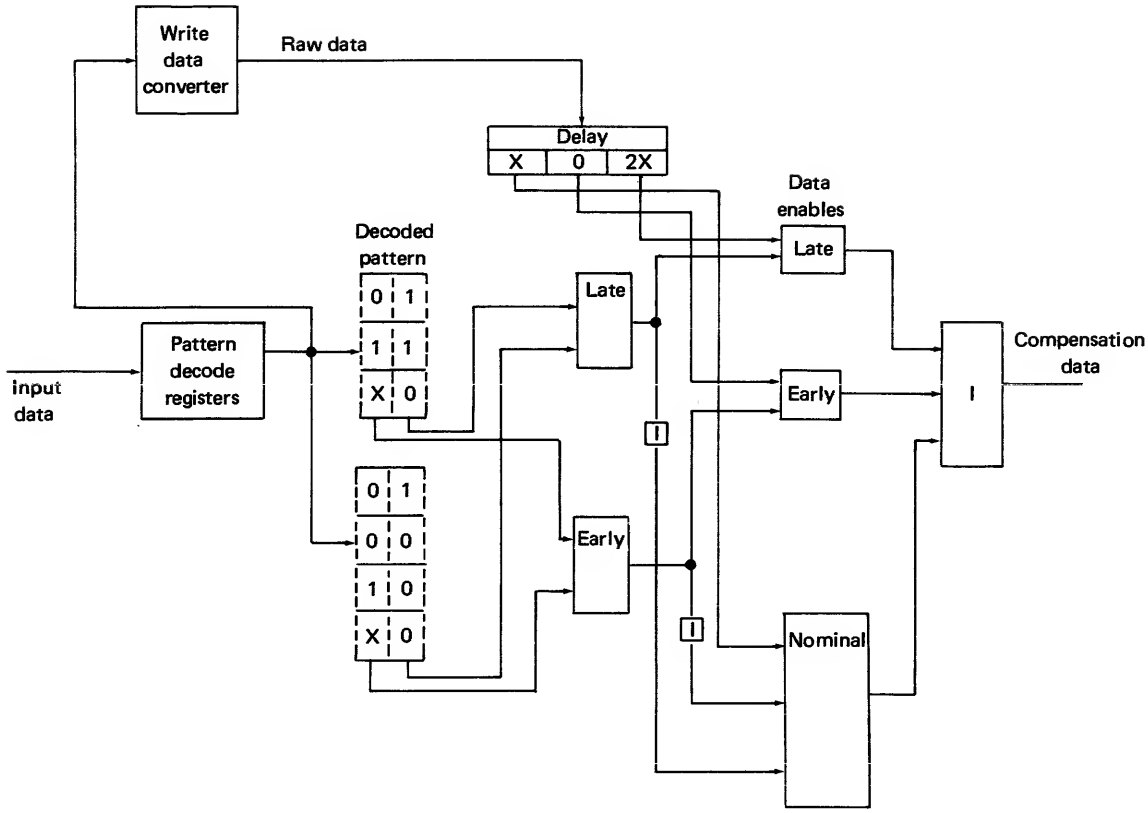

Figure 9-36. Write compensation block diagram

The circuit operation begins at the pattern decode box which receives the input data from the controller. Enable signals shift the data through the register to the write data converter and the pattern decode logic.

The write data converter converts the input data into the write data and then applies the write data to a delay line. The delay line has an output at the early, late, and nominal gates. These three outputs are combined with the outputs of the pattern decode logic to produce the compensated write data.

The pattern decode logic is formed by a series of gates that analyze the input data as it shifts through the pattern decode register. The input frequency is determined to be constant (0000 or 1111), increasing (011 or 1000), or decreasing (10 or 001). The outputs of the decode logic enable the early, late or nominal gates to provide compensated write data. This determination is made by the following method:

- If frequency is constant, there is no peak shift. Data is defined as nominal and is delayed by X.
- If frequency is decreasing, the readback signal peak occurs later than nominal. To compensate for this, data is not delayed.
- If frequency is increasing, the readback signal peak occurs earlier than nominal. To compensate for this, data is delayed by 2X.
- If frequency is constantly low, delays through the read chain cause the peak to come late, therefore data is not delayed.

After being write compensated, the data is transmitted to the write driver circuits.

The write compensation timing diagram (see figure 9-37) indicates how the outputs of each box may appear during a write compensation. A brief explanation of this follows:

- Input Data – data as arrived from controller.
- Raw Data – data as it appears after being shifted through pattern decode registers and converted.
- Early Pattern – The 001 pattern of the input data indicates frequency decreasing, sets early gate, and one side of the early data enable. Since only an X delay from nominal is desired, there is no delay of the write signal. Early data enable gate is set.
- Late Pattern – The 011 pattern of the input data indicates frequency increasing, sets late gate, and one side of the late data enable. Since an X delay from nominal is desired, the signal must be delayed by 2X to obtain this.

This compensation string continues for the entire line of data.

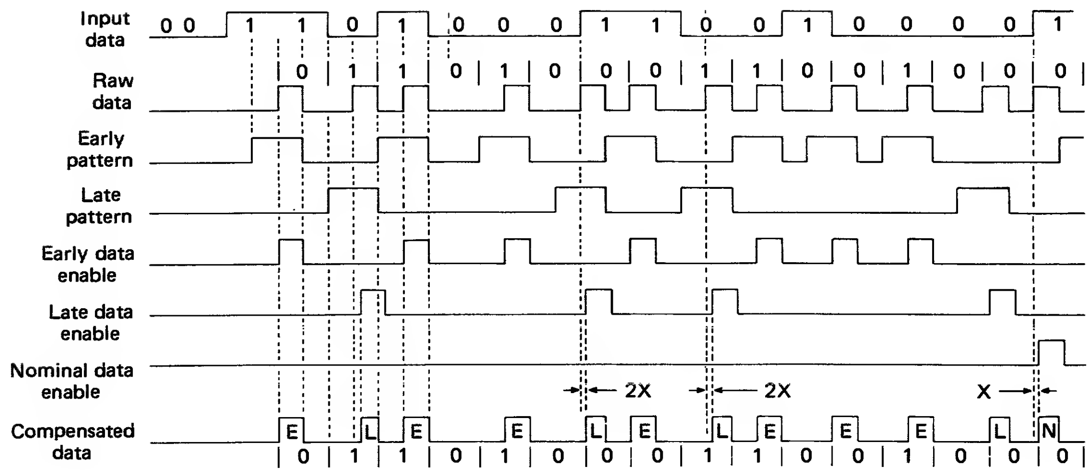

Figure 9-37. Write compensation timing

Current Zoning

The method used to compensate the write signal versus the effective linear speed on the disk is called current zoning.

The compensated data is sent to the write driver circuits which generate the write current for the heads. The magnitude of this current is controlled by the cylinder (track) address. Write current is maximum at the outer cylinders, and is reduced as each selected zone boundary is crossed. Therefore, as the heads fly closer to the surface, less current is required to saturate the medium.

The number of zones may vary from two zones to about eight zones, depending on the manufacturer. An example of zone selection is as follows: if there is a 200 TPI disk and the recording surface is four inches, the device is capable of recording 800 tracks (cylinders) of data. The current zones may be:

Zone 1	tracks 000 to 100
Zone 2	tracks 101 to 200
Zone 3	tracks 201 to 300
.	. . . .
.	. . . .
.	. . . .
Zone 8	tracks 701 to 800

As previously stated, the current zoning is controlled by the cylinder, or track, address.

Read Data

Read operations are initiated by the controller. This enables the analog detection circuits, which sense the data written on the disk and generate analog read data signals.

The analog read data is sent to an analog-to-digital converter which converts the analog signals into digital signals for processing. See figure 9-38 for a simplified block diagram of the read circuit. Later activities will explain the functions of these blocks in more detail.

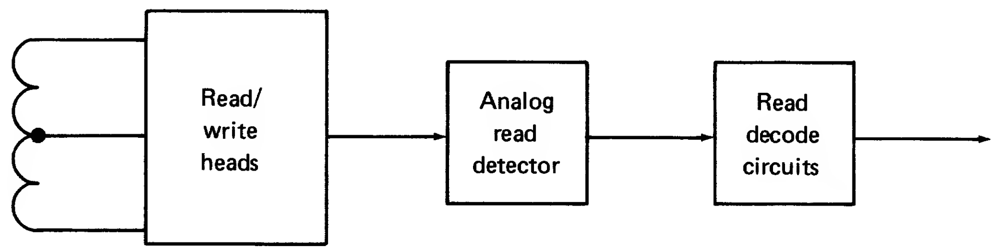

Figure 9-38. Simplified block diagram of read circuits

The induced voltage of the read/write heads is a few microvolts. The frequency of this analog voltage is proportional to the frequency of the magnetic field flux transitions sensed by the read coil. The voltage is transferred from the read/write heads to the analog read detection circuits. Here, the signals are amplified and filtered. Amplification is necessary to obtain a workable signal level; filtration is necessary to eliminate unwanted noise. The filtering attenuates the high-frequency noise and provides a linear phase response over the range of the read data frequencies.

The read decode circuits receive the analog read data from the analog read data detection circuit and convert it to digital data.

Introduction to Addressing

A disk address is used to locate data in much the same way that a home address is used to locate a person.

Addressing is required because of the vast storage capacity of a disk pack, and the fact that, often, a system writes short blocks of data it may have to read frequently. It can do this if the pack is divided into small addressable areas. However, it is also possible to record large blocks of data by using these small addressable areas in succession.

Usually, the system software is responsible for selecting the proper disk address for data. To write data on the disk, the system software must keep track of which addresses are not being used. It must also keep track of where data is stored, so that it can be retrieved at a later time.

A track, you remember, is the recording surface beneath one head at any given time. A track is one concentric circle on a disk surface (see figure 9-39).

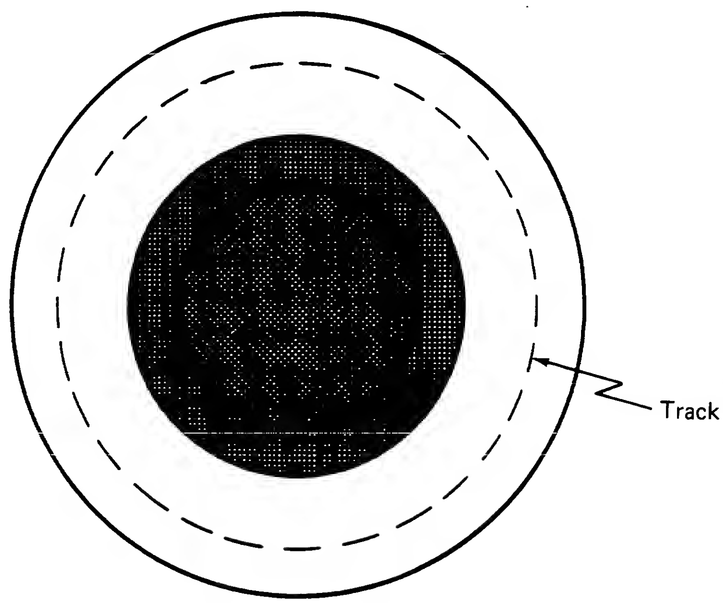

Figure 9-39. Disk surface

Cylinder is defined as the recording surface beneath all heads at any given time, or the total amount of recording surface available without moving the heads. Figure 9-40 shows how a cylinder is located on a typical disk pack.

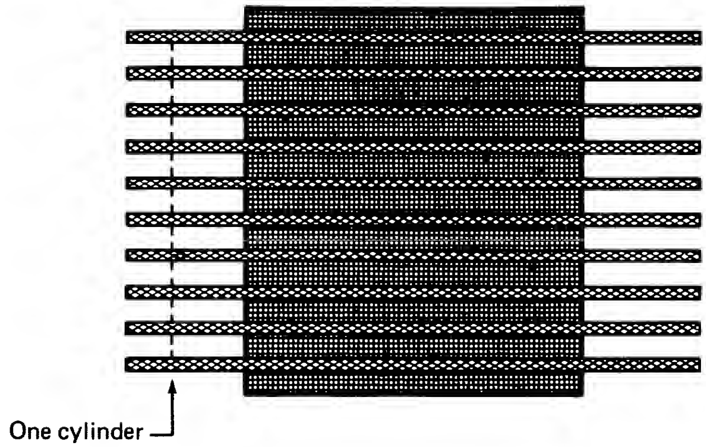

Figure 9-40. A cylinder

The first level of addressing is selecting the proper cylinder. The heads can be positioned to any recordable cylinder by the positioning system. Once the heads are positioned at the proper cylinder, the desired recording surface must be selected. You have learned how the desired head is selected. Selection of a head selects the recording surface that is to be used.

The final step in addressing on the disk surface is selecting the proper segment, or sector, of a track. Remember, the system software must know which cylinder, head, and segment (sector) to select. The software then causes the proper signals to be transmitted to the disk drive. The disk drive decodes these signals, positions the heads, selects the proper head, and selects the desired sector. You should note that every address on the disk pack is unique; there is no duplication of addresses.

To understand this final level of addressing, look at the two views of a disk surface shown in figure 9-41.

The disk on the left represents a typical disk surface and has the recordable area indicated. Note that the disk on the right is divided into eight pie-shaped areas. Each of these pie-shaped areas is called a sector. A sector is defined as a short segment of a track. Each of the eight sectors in figure 9-41 is exactly the same size as the other sectors, and is numbered sequentially from a point called the index. The sectors are electronically separated, just like the index is electronically formed. Index is a reference point, and is used to synchronize disk timing.

A disk having a single index marking is generally called a soft sector disk device. When this single marking is sensed, a timing sequence of pulses is generated to produce imaginary sectors for the remainder of the track. These imaginary sectors are called soft sectors.

A hard sector device contains mechanical markings on the disk surface to physically mark the positions of all sectors. The number of markings varies from one to thirty or more.

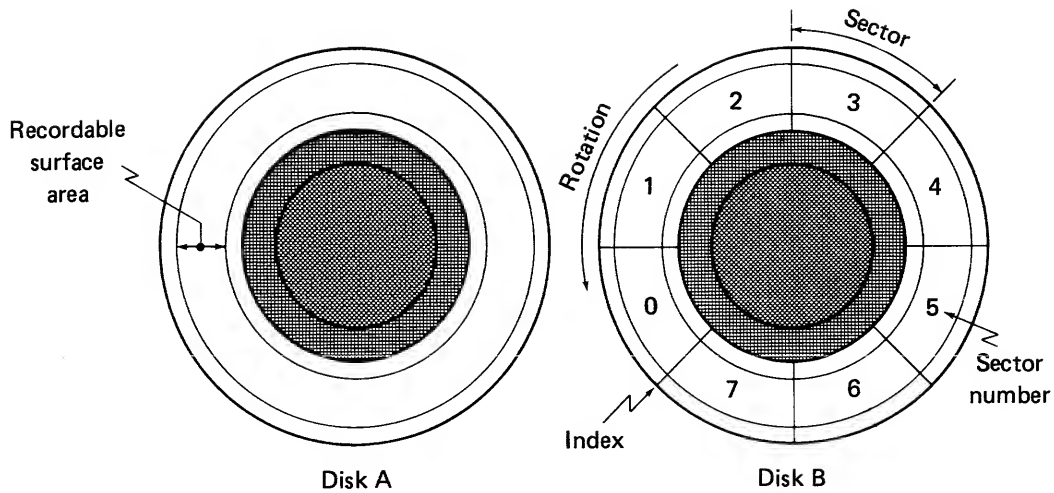

Figure 9-41. Surface segments (sectors)

Addressing a very small area on the disk is possible when the disk is divided into sectors. This allows very rapid location and retrieval of data from the disk surface, which is one of the primary objectives of disk storage.

If you examined a sector position in a typical track, it would be the same as illustrated in figure 9-42. Remember, this sector is exactly the same as all the other sectors on this surface and in this disk pack. If there are 400 tracks on this surface, with eight sectors per track, there are 3200 (400 x 8) individually addressable spots (sectors) in this surface alone. If the same surface was used as a 29-sector surface, the total number of sectors per surface would be 11,600 (400 x 29). Each sector would still have the same format as all other 11,599 sectors, and would contain similar data. The design of the disk drive logic circuitry controls the number of sectors on the surface of the disk.

Because all the sectors are the same, you can examine any sector. For example, look at sector 3 of the center track, which is highlighted in figure 9-42. If this surface is used as an 800-track surface, this track would be track 400, because it is midway in the recording area.

There are a number of fields within each sector. Each sector has the same fields, and they are located in the same sequence.

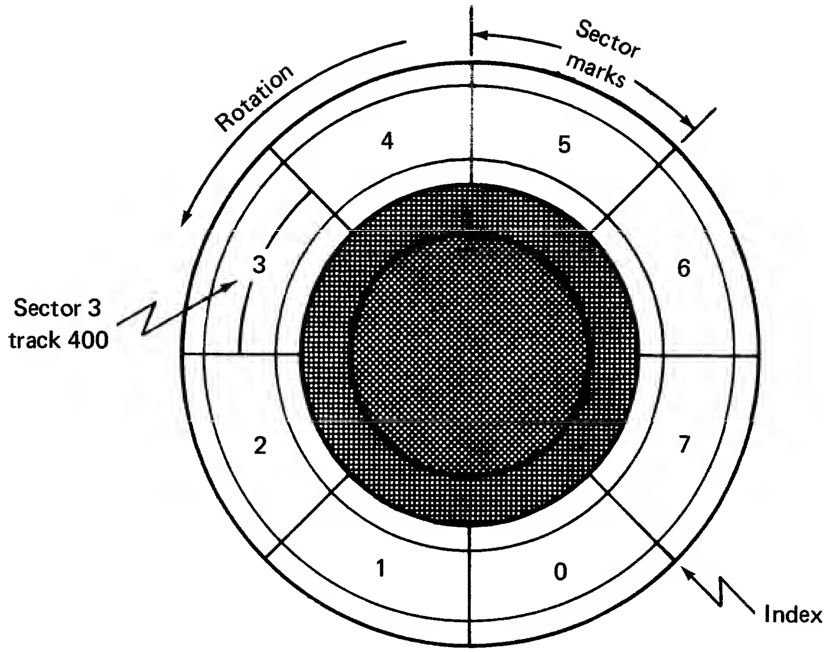

Figure 9-42. Sectors

Figure 9-43 shows the fields in a typical sector. Some disk drives may have more or fewer fields, but all the sectors on a disk pack are exactly the same size and contain the same fields. The physical size of a sector in an outside track is larger than the physical size of a sector in an inside track. However, each contains the same number of bits. These bits are just packed a little closer together on the inner track.

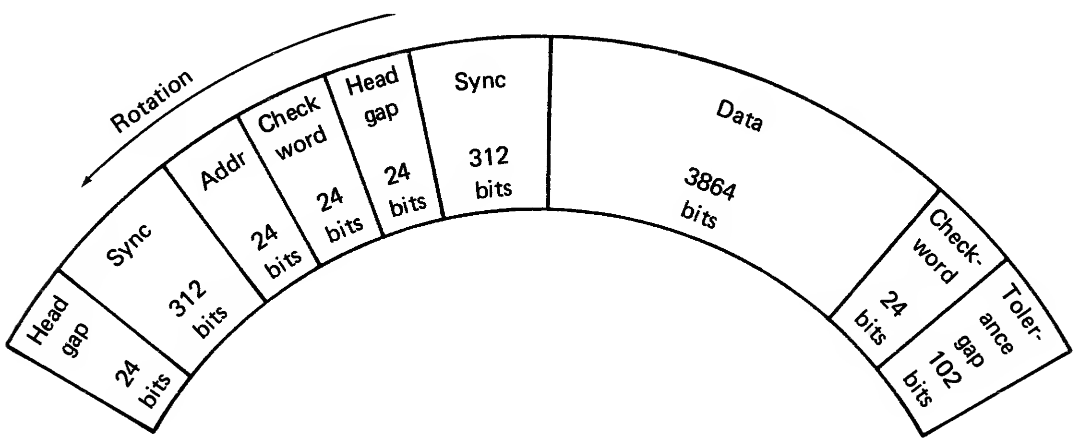

Figure 9-43. Sector fields

Each field within a sector has a distinct function in achieving the overall objective. Figure 9-43 shows the sequence in which the fields rotate beneath the head. These fields pass beneath the head from left to right.

The number of bit positions in each field is a feature of design. However, a field will be kept at the number of bits assigned to it. For example, if an address field is assigned 24 bits, it will always contain 24 bits, and if the sync field is assigned 312 bits, it will always contain 312 bits. (The only exception is the tolerance gap field, and the reason for this will be explained later.)

Head Gap

The head gap is an area where all current is removed from the read/write head and its associated circuits. This allows any circuit noise or stray magnetism to settle down. This is also important if the head function is switching from writing to reading.

Sync

The sync field is used to synchronize the speed of the disk to the clock timing of the disk drive. The field usually consists of a string of zeros with a special pattern at the end. The head is turned on at some point in the sync field and uses the string of zeros to begin the synchronization. Therefore, when the special pattern is read, it finalizes the synchronization, and, more importantly, signals the logic that the address field is coming beneath the head.

Address

The address of each sector is contained in the address field. As you can see in figure 9-44, the address field has at least three components. They are the cylinder, head (surface), and sector. An address can be pinpointed precisely by using these three components. Notice that, in keeping with the example, the cylinder number is 400, the head (surface) number is 6, and the sector of track 400 is 3.

The addresses are written at the factory, or by selected site personnel. A special software routine must be performed to write the addresses on the disk pack. The disk drive circuitry is designed to prevent writing in this area during data transfers; but occasionally (due to malfunction or operator or programmer error) writing will occur in the address area. This means that the disk pack must go through the special address writing routine to restore the addresses. The process of writing addresses is called formatting, writing address tags, or writing headers.

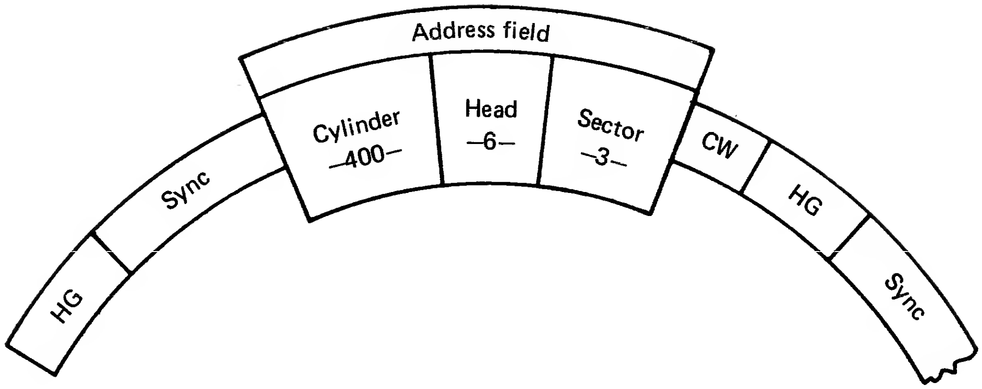

Figure 9-44. The address field

Checkword

The checkword is a specially coded word used for error detection. The checkword field in figure 9-44 is used for detecting errors in the information written in the address field. As each bit of data is written into the address field, it also passes through a logic circuit called a checkword generator. This generator uses the data bits to develop a special code. When the last bit has gone through, the special code is written into the checkword field. This checkword will be compared when the address is read, and an error check will be made.

The head will be shut off after this field is read.

Figure 9-45 has expanded the fields given so far. You should recall that they have all been associated with the address portion of a sector. The next four fields are associated with the data portion of a sector. Refer to figure 9-46 to see the position of each of these fields within the sector.

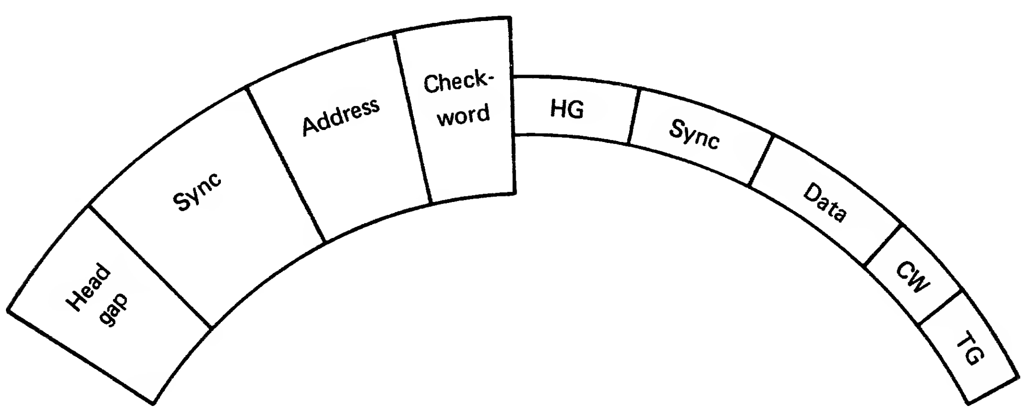

Figure 9-45. Sector fields (address)

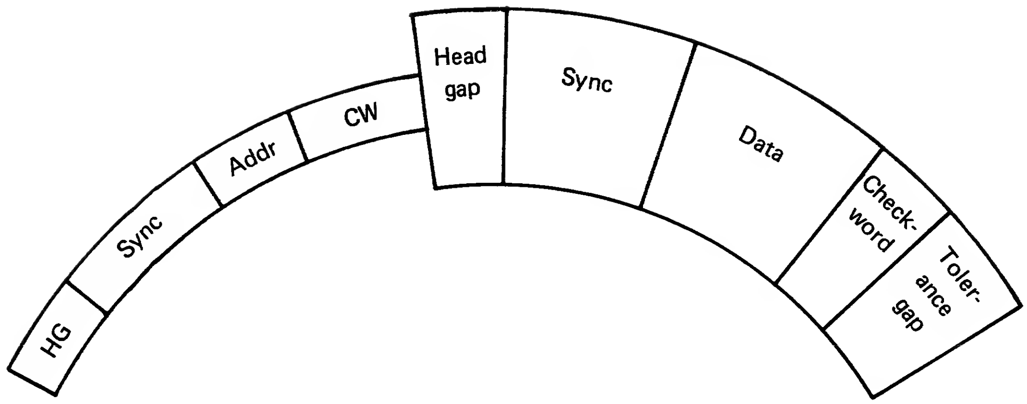

Figure 9-46. Sector fields (data)

Head Gap

This head gap is used in preparation for the data field in the same way the first head gap was used in preparation for the address field.

Sync

Because the head was shut off during the head gap, you must again use this sync field to synchronize the speed of the disk with the clock timing of the disk drive. This is necessary if you want to use the data field that follows.

Data

The data to be stored is recorded here. Each data word is sent to the head one bit at a time, and recorded on the surface one bit behind another. This serial mode of recording is the same for each word, because one word follows directly behind another. Refer to figure 9-47 for the data recording format.

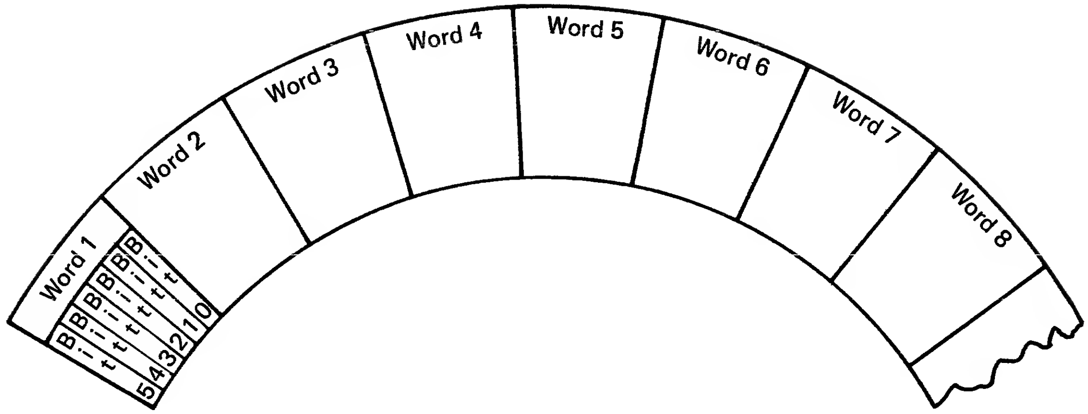

Figure 9-47. The data field format

Checkword

This checkword is developed by the data bits, which are recorded in the data field. It is used by the data field for error detection whenever that field is read.

Tolerance Gap

The tolerance gap (TG) is designed to compensate for variations in the speed of disk rotation. It ensures that there is enough time at the end of a sector to turn off the heads before the next sector mark appears. This prevents overlapping sectors. If one sector were to overlap another, both would be useless.

As you study the following two examples of typical disk addressing, refer back to figure 9-43 as needed, to understand the addressing operation.

Example 1

Conditions: The disk drive has just completed a return to zero seek, the addresses have been written on the pack, and the software in the computer decides to record a block of data on this disk at cylinder 126, surface 8, and sector 24. Assume that your pack has 800 cylinders, 20 recording heads, and 29 sectors.

1. The disk drive is commanded to seek to cylinder 126. (No data can be recorded while the heads are moving, so very little is going on while the seek takes place.)
2. The disk drive is commanded to select head 8 so surface 8 may be utilized. (This selection could occur while the heads are moving.)
3. The heads stop moving when they are ON-CYLINDER at cylinder 126.

4. When the ON-CYLINDER condition is developed, the circuitry waits for the next sector mark to appear, which means the beginning of a sector. (At this time, you have no idea which sector will be beneath head 8.)
5. As soon as the sector mark occurs, the head is turned off (it was probably off already), and the head gap field is rotated past the head.
6. The disk drive logic determines the length of the head gap, and, when the sync field is under the head, turns on head 8 at a predetermined point.
7. Head 8 begins reading a string of zeros. The clock pulses at the cell times will be sent to the disk drive logic. (These clock pulses will synchronize the speed of the disk with the disk drive read logic.)
8. When the sync pattern is read, the disk speed and the block of the disk drive logic will be fully synchronized. (This pattern will also alert the circuitry that the address field is about to be read.)
9. The address field is read and compared to the desired address. As each bit of the address is read, it is also sent to the checkword generator, where it develops a specially coded word as it did when the address was written. This specially coded word is then compared to the checkword that is read from the disk. This checks to see if the address data has been read correctly.
10. If the address field matches the desired address, and there is no checkword error, preparations are made to allow the head to be used in the rest of this sector. However, if the addresses do not match, the head is shut off until the next sector mark. If a checkword error exists, external circuitry is notified, and the addressing operation ceases.
11. After the address checkword of the desired sector is read, the head is shut off until the disk drive logic determines that the head gap is finished. This is done with special clock circuitry in the disk drive logic.
12. At the end of the head gap, the disk drive write circuitry turns on the head to write a predetermined number of zeros and a sync bit or pattern of bits. This writing includes the sync field.
13. As soon as the sync field is finished, the data field is written. Data writing continues until either the data block is finished or the data field has recorded its maximum number of bits. If the data block is shorter than the data field, the disk drive circuitry fills the remaining portion of the data field with zeros. However, if the data block is longer than the data field, the writing stops temporarily, when the end of the data field is reached, and begins again when the data field of the next sector is under the head.
14. During the writing of data, each data bit also goes to the checkword generator, where it helps develop a special code. This code is so intricate that a change in one bit of data may change several bits in the final checkword. When the last data bit has been written in the data field, the coded word in the checkword generator is sent to the head and recorded on the disk in the checkword field.
15. The head will be in the tolerance gap when the writing of the checkword field is completed. The head will be shut off and everything will be at rest until the next sector mark comes along. The length of the tolerance gap varies according to the speed of the disk pack. The tolerance gap allows for speed variations without affecting the size of the usable portion of the sector.

This was an example of writing in a specific sector on the disk surface. When the writing is finished, nothing more occurs until the next operation is requested by the system software.

Example 2

Conditions: Heads are positioned at cylinder 126, as they were in example 1. No head or sector is selected. The system software decides it needs the data recorded at cylinder 783, surface 17, sector 11.

1. The disk drive is commanded to seek to cylinder 783, and head 17 is selected.
2. When ON-CYLINDER occurs, indicating that the seek to cylinder 783 has been completed, the next sector mark causes a search for sector 11. This is done the same way it is done for writing. Every address in the selected track is read. If it does not match the requested address, the head is turned off until the start of the next sector. (When you locate the correct sector, you can continue with this operation.)
3. When sector 11 is present, the address and address checkword are read and verified. The head is shut off for the head gap, and will be turned on during the data sync field. The data sync field will synchronize the speed of the disk to the clock timing of the disk drive.
4. When the sync bit or pattern is read, it will flag the logic that the data field is ready to be read. The data field is read serially, bit-by-bit, and word-by-word. The words are sent out to the external circuitry. If the amount of data requested is less than the amount of data in the data field, data transmission to the computer stops when the required number of words have been read. The disk drive continues to read to the end of the field, and reads and checks the checkword. If more data is required than one data field has available, the reading of data will temporarily stop at the end of this data field and continue reading when the data field of the next sector is under the head.

Error Control and Correction

This activity lists some of the causes that may create errors in magnetic storage devices. You will look at a cause and determine if it could be internal or external to the device and then determine what possible effects it has on the equipment. You will also look at possible methods to control these causes and at methods used to correct errors that exist.

Causes of Errors in Magnetic Storage Devices

There are many different causes of a read/write error in magnetic storage devices. These errors may be externally or internally induced into the system.

Externally induced read/write errors may be caused by one or more of the following conditions:

- Improper or insufficient site grounding system
- Induced noise from other equipment
- Input power fluctuations
- Improper or insufficient air conditioning
- Excessive dust and dirt

Internally induced read/write errors may be caused by one or more of the following conditions:

- Improper or insufficient equipment grounding
- Improper handling of recording media and/or equipment
- Improper or insufficient preventive maintenance
- Poor quality recording media

Though this section is dealing with magnetic storage devices, the causes of failures noted above apply to all computer and peripheral equipment.

Grounding

There are basically three types of ground systems for electronic equipment:

- Signal ground – The return lines for both input and output power; this is also known as the zero reference level.
- Chassis ground – Used to ground the equipment chassis and frame to a zero potential. This eliminates static charges from building up on the frame and prevents personal injury from static charges, improper AC connection, and even some internal shorts.
- Earth ground – The single point for all equipment grounding. As the name implies, this ground is the earth.

Sufficient grounding is necessary to reduce or eliminate noise, ripple, or other spurious responses which may be induced into the circuits.

Proper grounding is necessary to avoid ground loops which would act as antennas and induce radio frequency or electromagnetic interference into the circuits.

Grounding problems could result in timing errors, false start/stop signals, or false read/write signals; any of these could be the cause of extra or missed data.

Induced Noise from Other Equipment

Relays, motors and devices of this type may have a reverse EMF, which induces a noise spike on the AC power line. Some line printers, typewriters and teletypes, when keyed, are the most probable cause of this spiking. This noise spike is induced into all equipment on that line.

Power Fluctuations

This is the fluctuation of input power due to overloading or heavy current drain. This is most prevalent when equipment is first energized or during periods of heavy electrical usage.

Air Conditioning

Temperature increase is one of the worst enemies of proper equipment operation. Temperature coefficients of materials affect the entire performance of every component (electrical and mechanical) in a piece of equipment. Heat can cause bearings to bind, slow rotation, alter the magnetic properties of material, alter flying heights of recording heads, alter signals, alter timing, stop circuits from working, and on and on.

Humidity is another concern – too much humidity can cause moisture which can create arcing in high voltage supplies, or even allow signal crosstalk. Too low of a humidity level allows the build-up of static charges, which introduce read errors.

Dust and Dirt

Dust and dirt can clog filters, block air flow, retain heat in components, and prevent cooling of circuits. Dirt can clog bearings and increase friction, thereby increasing heat. Dirt can coat recording surfaces and heads, and cause a head to crash.

When moisture from humidity combines with dust, the dust acts as a capacitor, allowing signal crosstalk or the generation of noise.

Equipment Handling

Every magnetic storage device (tape, disk or memory), has conductors to carry the data. These conductors are routed in a manner that reduces the possibility of induced cross-talk because of magnetic or electric field coupling. Care must be exercised when working around these conductors so that this effect is not created. Do not change or reroute these conductors.

Error Control

Obviously the best method of error control is complete and correct equipment/media handling and preventive maintenance. Several other considerations for error control are:

- All heads on all disk drives, using the same interchangeable packs, must have the same alignment in relation to actuator and disk spindle. This will help prevent peak shift and other offset problems due to skew.
- A brief temperature stabilization time should be allowed for cold disk pack/warm machine. This will help the possible misregistration due to thermal expansion.

As can be seen in figure 9-48 the room air is directed or channeled to four distinct sections of the disk drive: both power supplies, the logic chassis, and the spindle/actuator area. The spindle/actuator area is also isolated from the heat generating sections of the drive.

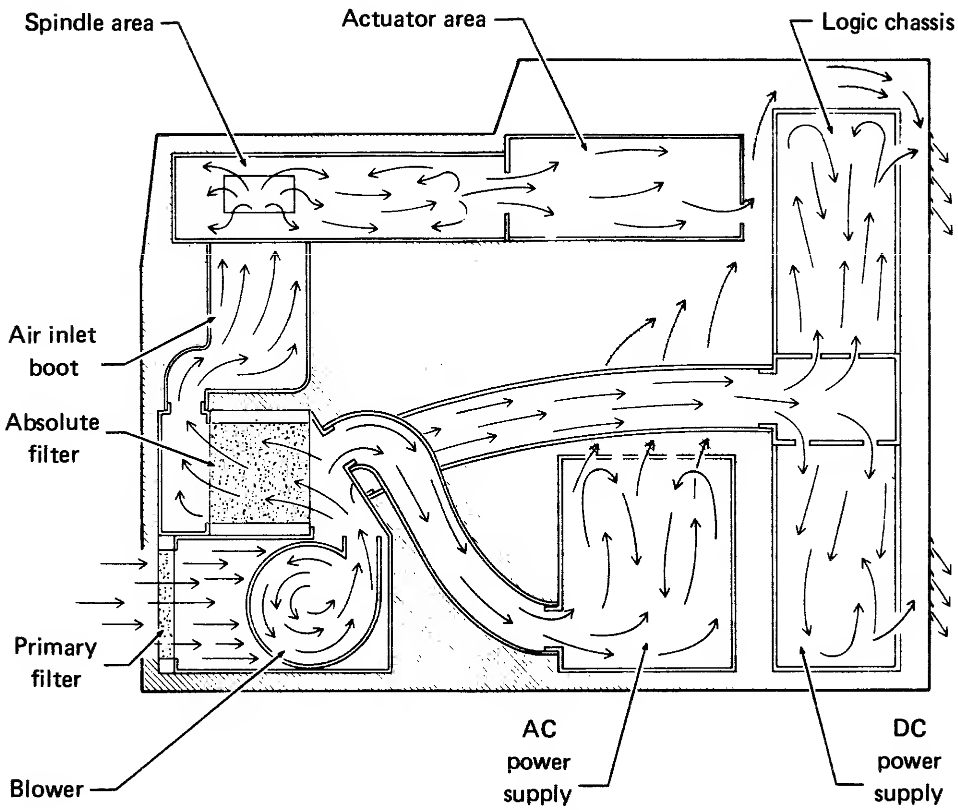

Figure 9-48. Blower system

The primary filter stops the larger pieces of dust and dirt from entering the drive. The absolute filter is a series of very small filters that prevent most dust from entering the spindle/actuator area.

The blower also assures a fairly constant air pressure in the spindle area and cools the heat generating areas. Though the major assemblies may vary in location from disk model to disk model, the same basic function of the blower system applies.

Equipment Grounding

Follow the manufacturers guidelines for correct and proper equipment grounding. These guidelines will usually be found in the reference manual supplied by the manufacturer.

Temperature/Humidity

These specifications will also be found in the reference manual supplied by the manufacturer. Though you may not have direct control over these items, it is your responsibility to inform the customer when these specifications are not being met.

Temperature will be your main concern, as the humidity level will be hard to determine.

Error Control and Media Defects

Every disk system has a combination of error correction techniques. Some of these techniques are:

- Checkword
- Skip displacement
- Sector/track deallocation
- Carraige offset
- Data strobe offset

These techniques may be hardware oriented, software oriented, or a combination of the two; the techniques vary among manufacturers.

Checkword

Address and data transfers are checked for accuracy by generation of a bit redundancy checkword in the controller. During operations the controller generates and verifies checkwords to determine the correctness of data transferred between the controller and the various peripheral storage devices.

The checkword is a cyclic code generated from the address or data being transferred. The checkword (remainder) is the bit code left in the cyclic encoder after the last bit of the address or data has entered the encoder. This is then written immediately following the address when the address is originally written or at the end of the sector for each write function. The remainder is obtained by using a mathematical expression called a polynomial. During subsequent address verification or read operations, the address or data being read is again fed into the encoder and a new checkword is generated. The checkword previously written then enters the encoder. If the original data was written and read correctly, the two checkwords cancel each other such that the encoder ends up in a clear state. (If any stage of the encoder is set upon completion of the operation, an error has occurred in either the original writing on the storage unit or during the subsequent read operation, and a checkword error indication is generated.)

Figure 9-49 shows a simplified block diagram of the cyclic, or checkword, encoder.

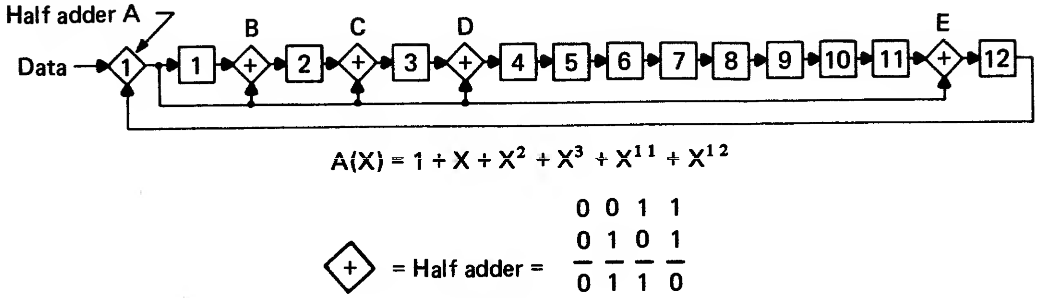

Figure 9-49. Block diagram of cyclic encoder

With the latest technologies, the cyclic encoder is able to determine which bits of data are in error and automatically correct them.

Skip Displacement

Skip displacement is one of the latest techniques in handling media defects. Usually, when a new disk pack is installed, a surface test is run on it. The surface test locates any defects which exist on the media. This surface test may also be called "mapping." Skip displacement is then a written header on the pack, which tells where these defects are. The defects are known by location and size. Thus, when the location is reached, data will not be written for X number of bytes. In other words–it skips over the bad spots.

Sector/Track Deallocation

This is the same as skip displacement except that entire sectors or tracks may be skipped.

Carriage Offset

It is virtually impossible to have the exact same head alignment on all disk drives in relation to actuator-to-spindle distance. Therefore, carriage offset is used to allow the head minute adjustments, when on track, to locate maximum signal strength.

Temperature also affects the track location area.

Data Strobe Offset

As with carriage offset, data strobe offset allows the data strobe timing to be adjusted to compensate for minute differences in rotational speed and other timing differences. Several passes are made at retrieving the data. If the data is still in error, either one or both of the offset techniques are used. The number of passes attempted for each technique may depend on hardware, software or both.

Introduction to Heads

What are read/write heads and how do they move? This activity introduces the components, assemblies, and movement of the heads.

Read/Write Assemblies

The reading and writing of data on a rotating magnetic storage device is accomplished by a device called the head/arm assembly. The head/arm assembly is also referred to as the head assembly or just the head.

The head/arm assembly is composed of the following three subassemblies:

- Fixed arm assembly
- Floating arm assembly
- Plug and cable assembly

Fixed Arm Assembly

The fixed arm assembly, also called the rigid arm assembly, is the mechanical support. It is connected to the carriage assembly of the actuator.

Floating Arm Assembly

The floating arm assembly is connected to the fixed arm. The floating arm assembly consists of three subassemblies:

- Read/write core
- Head pad
- Gimbal spring

Plug and Cable Assembly

The plug and cable assembly is the electrical connection onto which all data is transferred. One end of the cable is connected to the read/write core, the other end is connected to the read electronics.

To summarize, each head/arm assembly is composed of the following:

- Fixed arm
- Floating arm
- Read/write core
- Head pad
- Gimble spring
- Plug and cable

See figure 9-50.

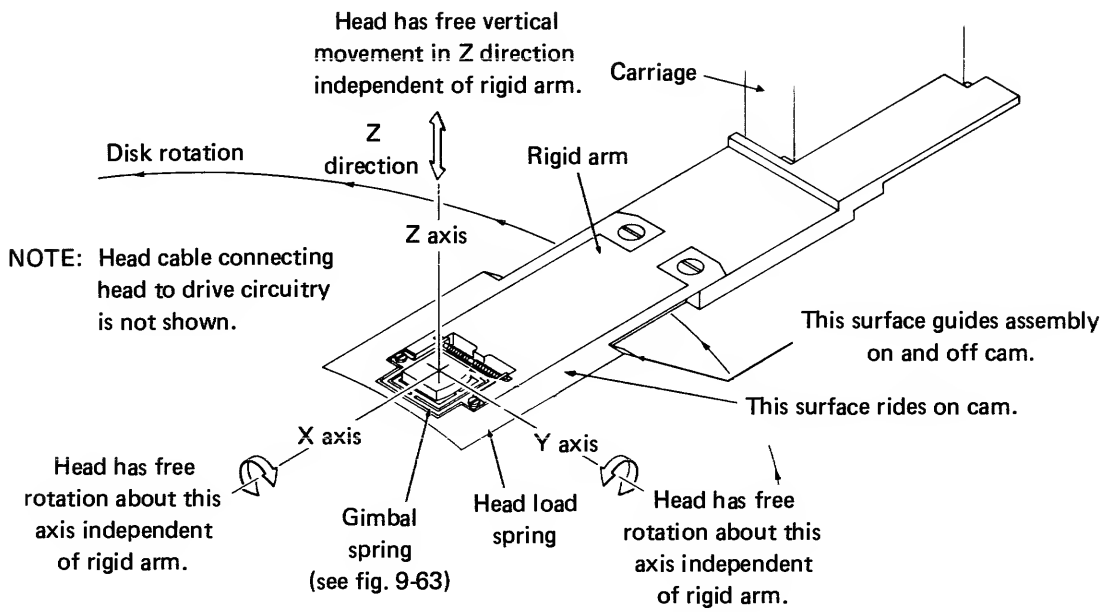

Figure 9-50. Head/arm assembly

Each of these components will be covered in greater detail in later activities.

Over the years, the head/arm assemblies have changed dramatically in their construction. Study figures 9-51 through 9-53; they show some of these changes.

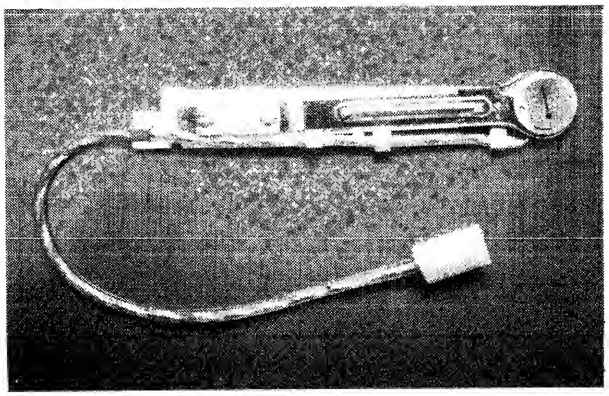

Front

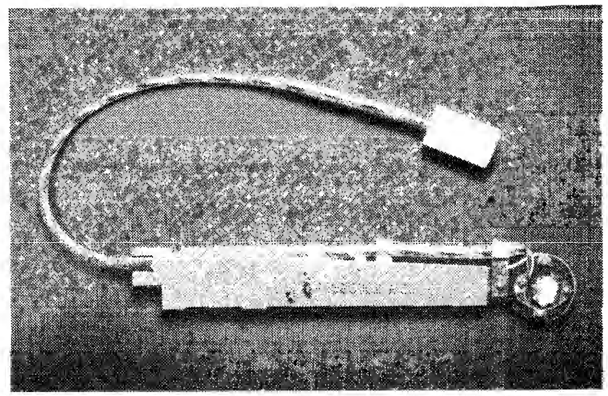

Back

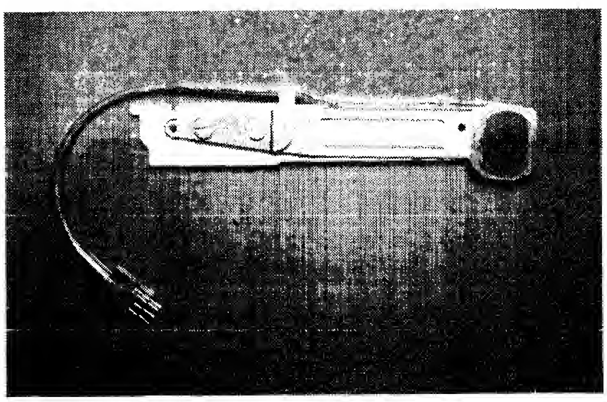

Front

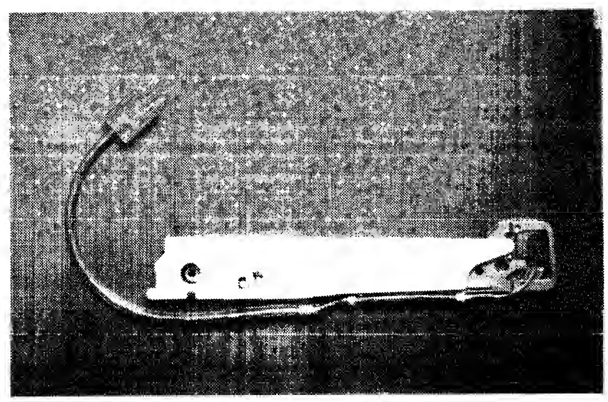

Back

Figure 9-51. Head/arm assemblies—circa 1966

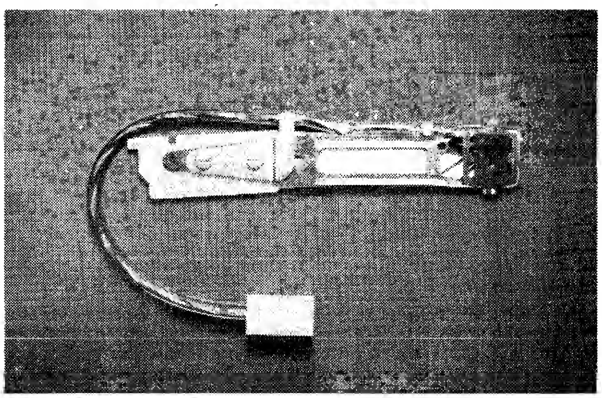

Front

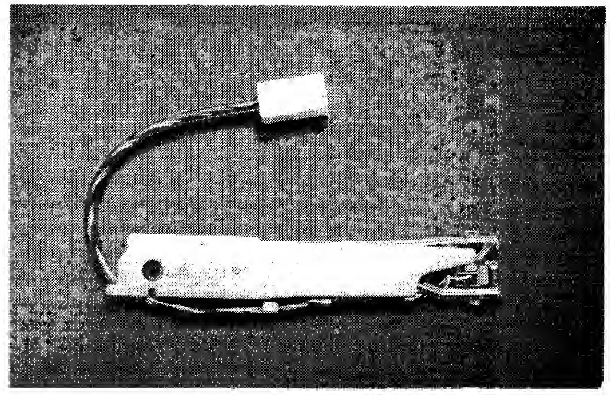

Back

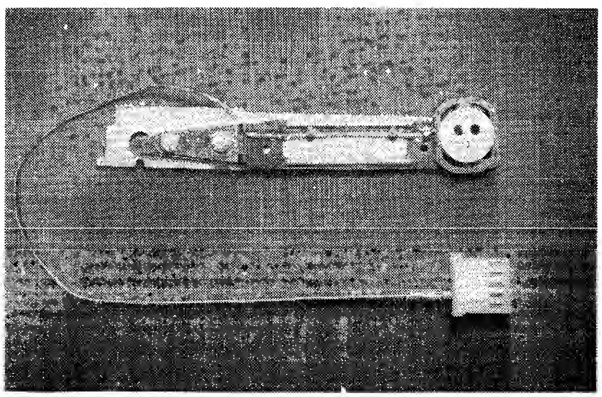

Front

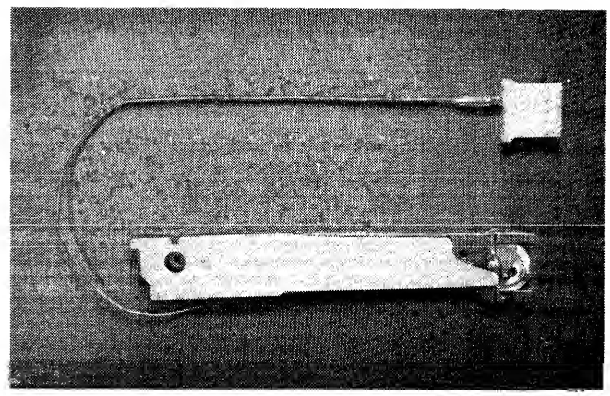

Back

Figure 9-52. Head/arm assemblies—circa 1970

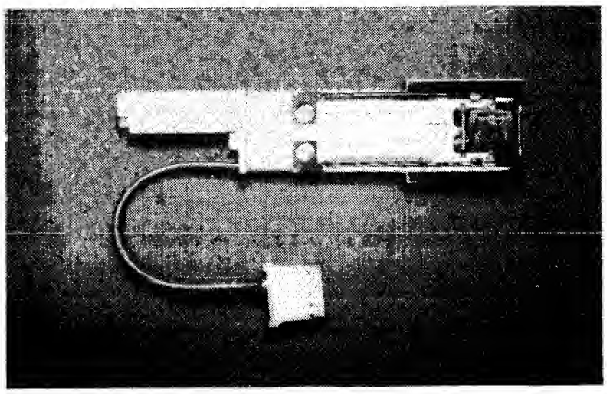

Front

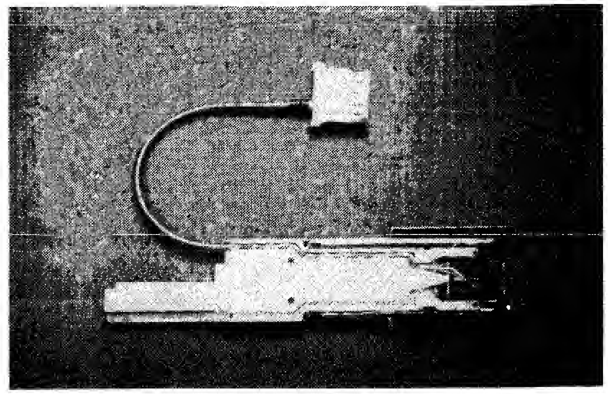

Back

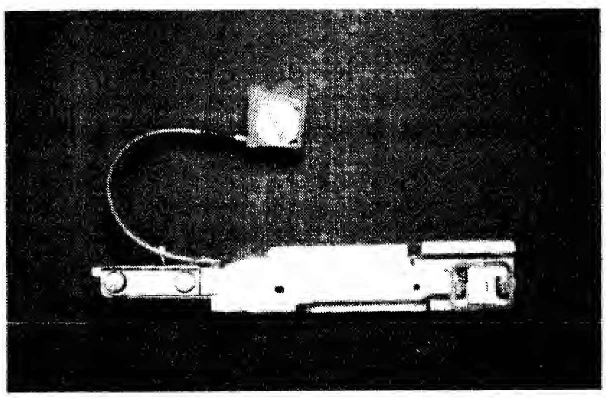

Front

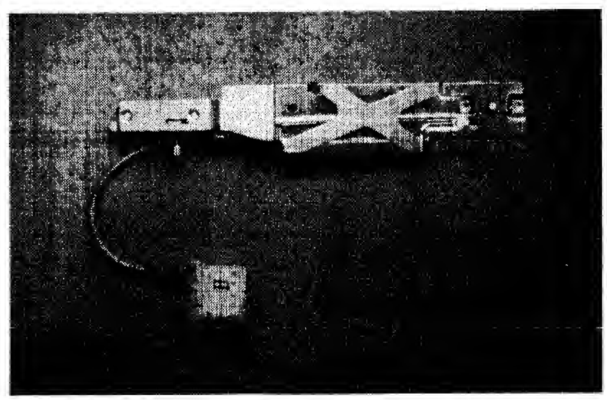

Back

Figure 9-53. Head/arm assemblies–circa 1974

Actuator

The actuator is the device that supports and moves the read/write heads. The actuator consists of the carriage, (on which the fixed arm is mounted), actuator housing, and a method to move the carriage. The method most commonly used today is voice coil. (See figure 9-54.)

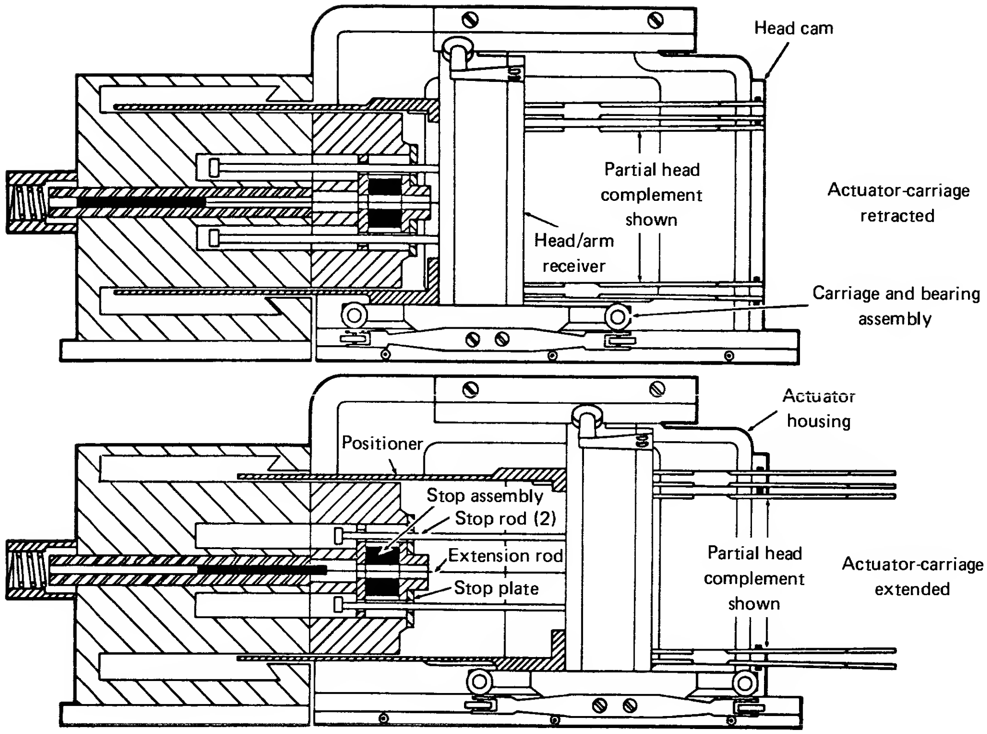

Figure 9-54. Actuator assembly elements

The forward and reverse movement of the carriage on the carriage track is controlled by a servo signal. The positioner is free to slide in and out of the actuator assembly. Fastened to the positioner is a head/arm receiver which holds the read/write heads. The head/arm receiver is mounted on the carriage and bearing assembly, which moves along the track on bearing-type rollers. Movement of the positioner, in or out, creates the same motion in the entire carriage assembly. This linear motion is the basis for positioning the read/write heads to a particular track of data on the disk pack.

Head Components

This activity discusses each component in a read/write head and explains how the components interact. The function of each component is described in detail.

The following components make up a read/write head:

- Magnetic core
- Head pads
- Gimbal spring
- Floating arm

Magnetic Core

The magnetic core is the smallest component in a head/arm assembly. The core is also one of the most important and secretive assemblies in a disk system. It is the device that writes or reads data on the disk media. Figure 9-55 shows some of the cores used in the past years; note the size differences of these cores.

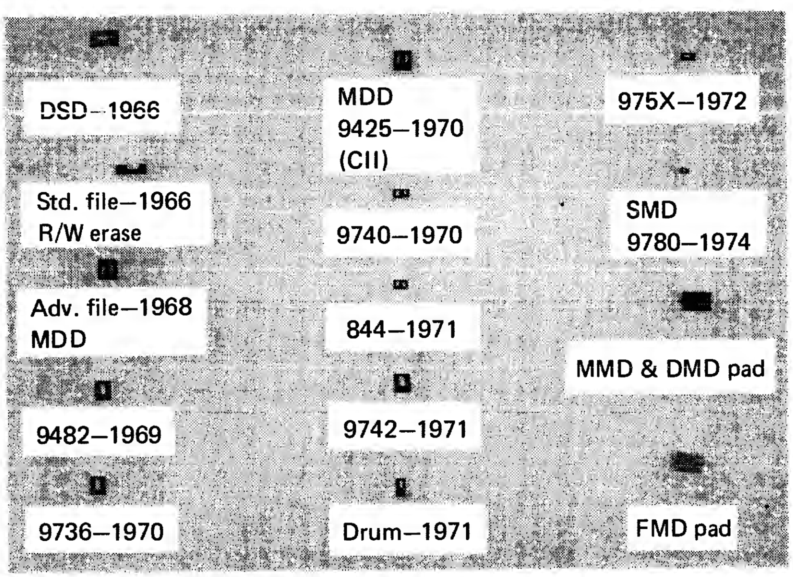

Figure 9-55. Magnetic cores

The magnetic core is a thin slice of ferromagnetic material which is slightly circular or horseshoe shaped. Several turns of 40-gauge, enamel coated wire are wrapped, bifilar fashion, around one side. This "coil" is center tapped thus forming three leads. The plug and cable assembly is connected to the three leads. Figure 9-56 shows one of the latest styles.

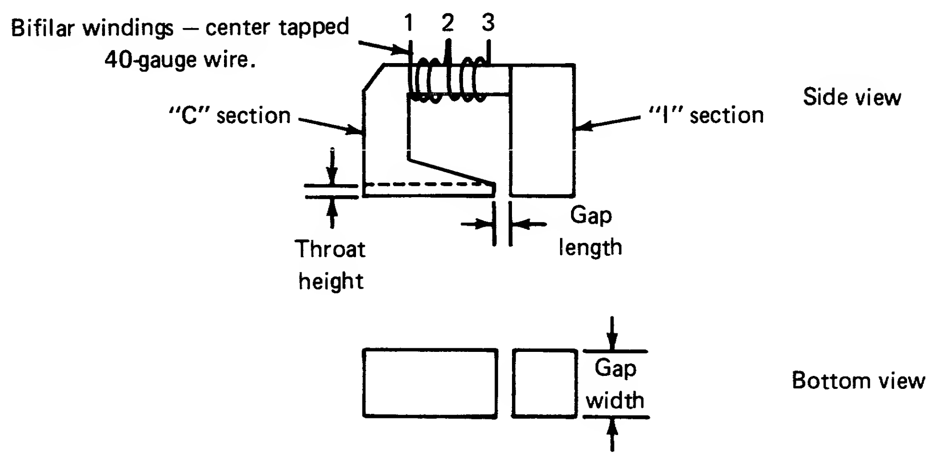

Figure 9-56. Read/write core

The combination of the coil and core form the transducer. When write current flows through the transducer, the coil generates magnetic flux. If enough current flows, the magnetic flux saturates the throat height. This flux leaks or spreads out at the gap, thereby causing the flux to go down into the media. This saturates the media. The flux then comes up through the "I" section and completes the magnetic circuit (see figure 9-9-57).

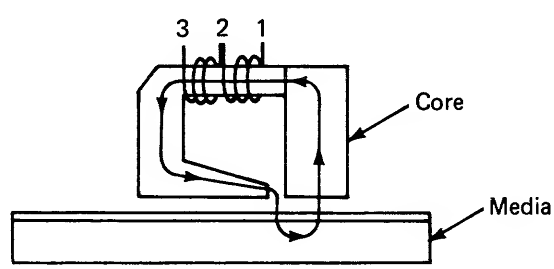

Figure 9-57. Saturation path

Throat Height

The throat height determines the amount of saturation current that is required in the coil. The thicker this gap is, the greater the write current required for a fixed number of turns. This increased current is necessary to force the lines of flux into the media. The nominal throat height dimensions are one microinch to several mils.

Gap Length

The gap length is one of the factors that determines the bit density of the disk drive. The narrower this length is, the higher the available bit density is. The nominal gap length dimensions range from 30 microinches to 110 microinches.

Gap Width

The gap width is one of the factors involved in determining the track density of the disk drive. The narrower the width is, the higher the available track density is. The nominal gap width dimensions range from 1 mil to 2.3 mils.

Head Pads

The completed core is securely mounted in the head pad. The head pad is made of non-conductive, nonmagnetic material. The function of this device is to position the R/W core at a fixed height from the media, without contacting the media.

As with the cores, the head pad has changed over the years. Figure 9-58 illustrates some of these changes.

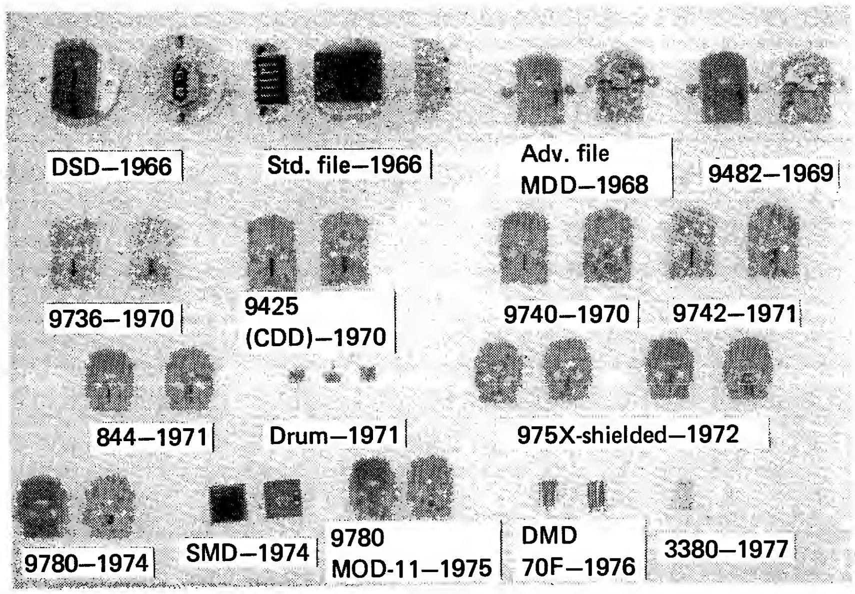

Figure 9-58. Head pads

Today, there are basically two types of head pads, these are called the crown pad and taper flat pad (see figure 9-59).

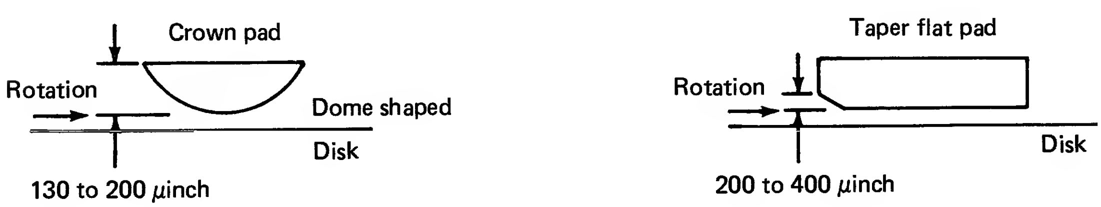

Figure 9-59. Head pads

Stabilization

The head pad acts as a dynamic air bearing. The head pad and gimbal spring allow the head core to pitch and roll as the disk requires.

Pitch is the up and down motion of the front and back edges of the pad. Roll is the up and down motion of the sides of the pad. Yaw is the horizontal rotational motion of the pad around its pivot point (see figure 9-60).

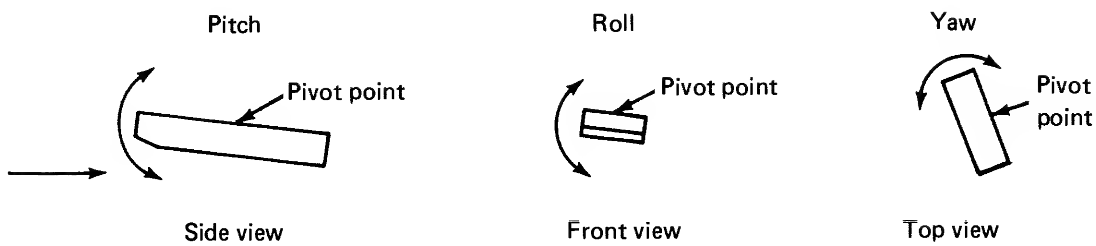

Figure 9-60. Basic pad motions

One method employed to lower the flying height of the head pad is the use of spoiler holes. Spoiler holes are holes placed in the head pad in several locations; the spoiler holes allow pressure release. These holes bleed off the air bearing pressure and lower the flying height of the head. However, as mentioned before, improper preventive maintenance allows dust and dirt to accumulate in these holes which decreases their efficiency. No attempt should be made to clean these holes without complete removal of the head/arm assembly, and then, only if that procedure is recommended by the disk manufacturer. The reason for this is that the dust becomes tightly packed and there is no room between the heads to vacuum or blow them clean and the use of a solvent would only loosen part of this dirt and smear the rest over the surface of the head pad/core.

Core Location in the Pad

The core should be as close as possible to the media. Therefore, the core is designed to be at the closest point between the pad and the disk surface.

In order to maintain maximum saturation of the media, the read/write core must maintain this close proximity to the disk. As the air pressure comes into the head pad, it tends to lift the pad (see the pitch view of figure 9-60). This places the back section of the taper flat pad closest to the disk. Therefore, the location of the core is in this area (see figure 9-61).

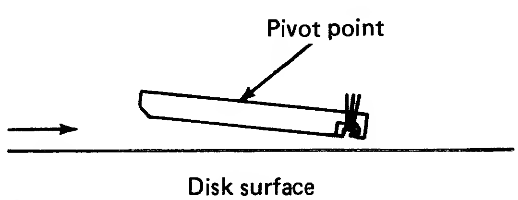

Figure 9-61. Core location–taper flat pad

However, in the crown pad, this placement becomes slightly more difficult. The core location on this pad is more dependent on the air pressure and rotational speed. As these increase, the tilt angle of the pad also increases; thus, the closest point can cover a wider area (see figure 9-62).

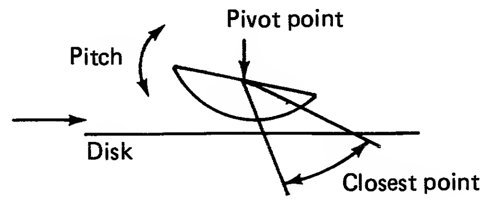

Figure 9-62. Core location–crown pad

During the design of a new disk system, the head pad air bearing is analyzed on a large-scale computer.

The computer calculates the pressure/surface variance of the head pad. This operation requires several hundred thousand memory locations.

Once the head core is assembled in the head pad, it is mounted on a gimbal spring.

Gimbal Spring

The gimbal spring is an extremely flexible piece of metal that attaches the head/pad assembly to the floating arm assembly (see figure 9-63).

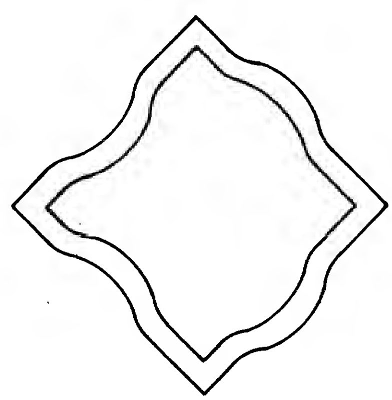

Figure 9-63. Gimbal spring

This allows the head/pad assembly flexibility of movement in two planes (pitch and roll).

Floating Arm

The floating arm is designed to maintain a constant loading force. This loading force acts in the opposite direction of the air pressure to prevent excessive disk/head clearance.

Suspension Systems

This activity introduces disk drive suspension systems.

To prevent destruction of either the disk pack or head/arm assemblies, the suspension system of the head/arm assembly must be carefully designed.

You can begin to understand the complexity of the design of the suspension system by considering the following points:

- The slightest contact with the media surface by the R/W head could destroy the media.
- The R/W head is flying above the surface of the media at a distance equal to 1/60 the thickness of a human hair.
- The speed of the media is in excess of 100 miles per hour.
- The head/arm assembly must be able to start from a dead stop, move several inches, stop, and process data within a matter of milliseconds.
- The head/arm assembly must be able to track the media surface regardless of the distortion due to these excess speeds.
- The head/arm assembly must be able to track the media surface regardless of any imperfections existing in or on the media.
- The combinations of movements by all head/media components could reach the resonant frequency of one of them. This would cause vibration, which could be catastrophic.

Suspension System

The suspension system in the disk drive consists of a floating arm and air bearing pressure. (Air bearing pressure is covered in another activity.)

The equivalent circuit of the suspension system is shown in figure 9-64.

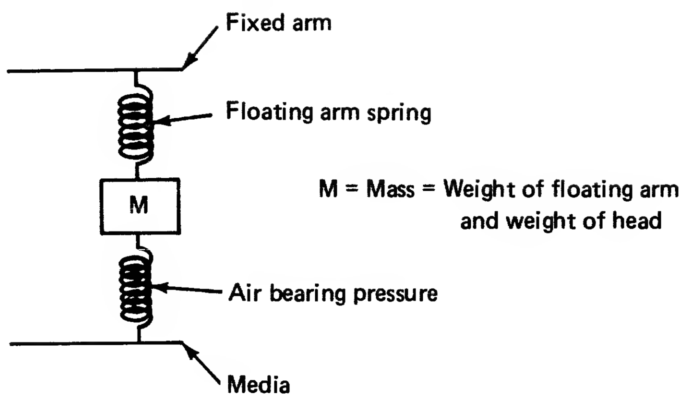

Figure 9-64. Equivalent circuit

An analogy of this is the car suspension system (see figure 9-65).

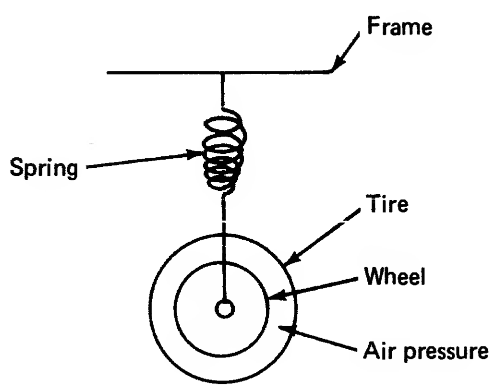

Figure 9-65. Analogous circuit

The intent and purpose of both of these circuits is to smooth out the ride. One circuit tracks the surface to avoid contact and destruction while maintaining accurate data processing. The other circuit tracks the surface of the road to remove obstructions to the ride, thereby making it more enjoyable.

Floating Arm

The floating arm is a very light material which is connected to the fixed arm by means of a flexible spring. The earlier disk systems used an arrangement in which the spring was attached to both the flexible and fixed arms (see figure 9-66).

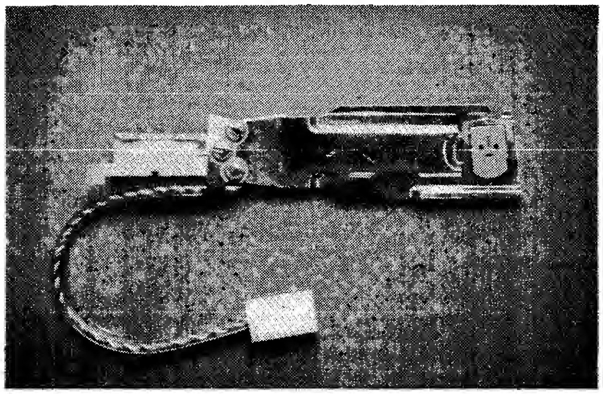

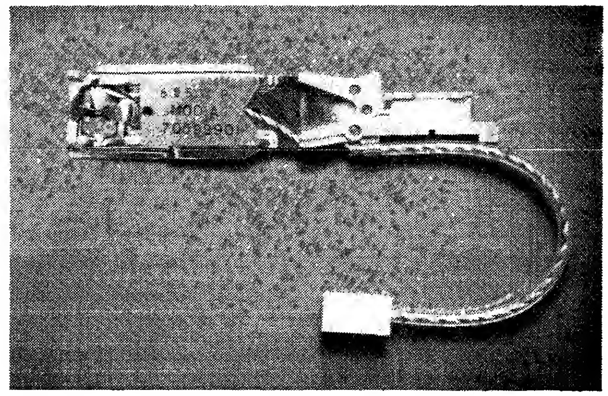

Figure 9-66. One-piece floating arm/spring

This spring may have been shaped in the floating arm as shown in figure 9-66 or it may have been a separate component as shown in figure 9-67.

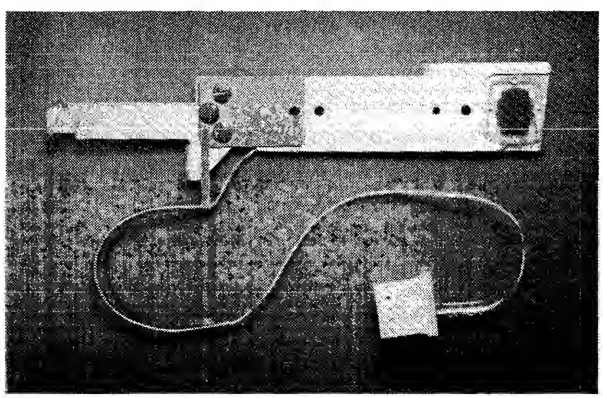

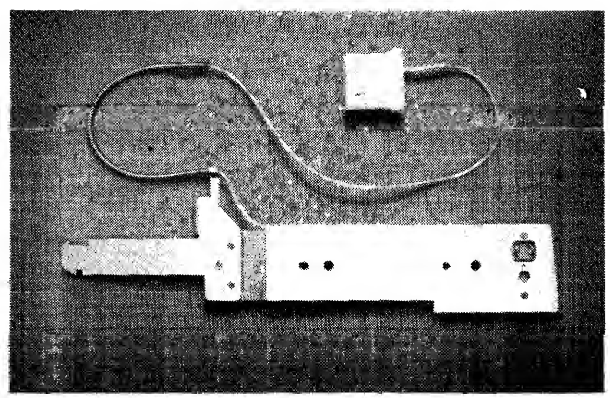

Figure 9-67. Two-piece floating arm/spring

As technology improved, so did the arm assemblies. Today, most of the mass of the flexible arm has been removed. The spring is a coil type mounted just behind the head/pad assembly (see figure 9-68).

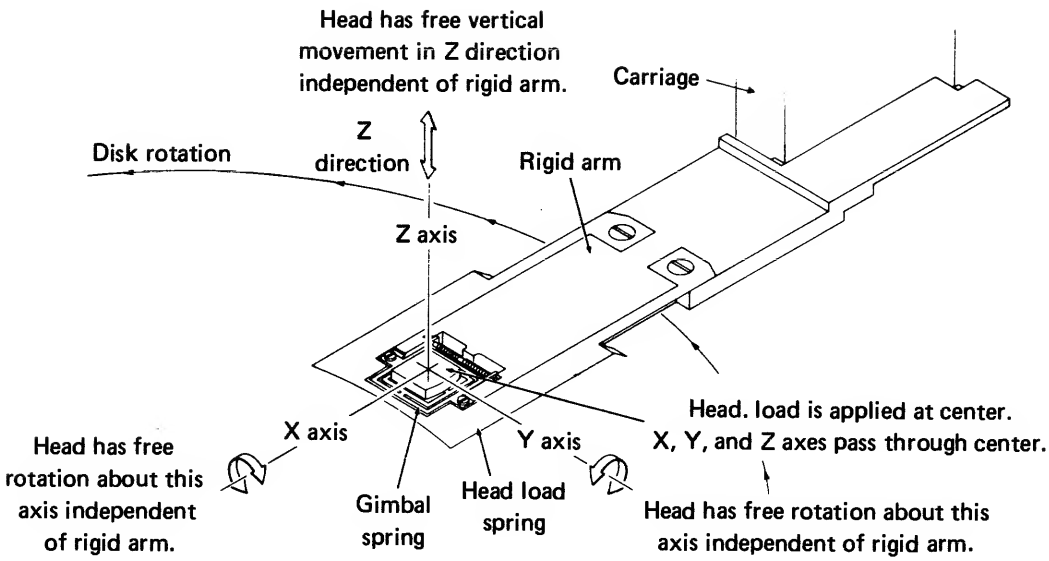

Figure 9-68. Present arm assemblies

Floating Arm Operation

As you can see in figure 9-68 the head/pad has three planes of movement:

- X axis
- Y axis
- Z axis

The major difference between the early assemblies and the present assemblies is that the total mass (M) is much lighter. This lighter mass, or weight, allows the head to fly closer to the disk surface because it has a faster response time to movement. This is particularly true of the Z axis.

If an extremely small imperfection existed in or on the media, on either side of the X axis, the head pad would respond in the X axis plane by means of the gimbal spring.

If an extremely small imperfection existed in or on the media, in the center of the head pad assembly, the head would respond in the Y axis plane by means of the gimbal spring.

If either of the above imperfections was slightly larger, or if the disk was not perfectly flat, then the head arm assembly would respond in the Z axis.

It is possible to have rotation about the x axis and the y axis and movement along the z axis all at the same time.

Resonant Frequency

All material has a natural frequency of oscillation. When this oscillation responds to an impulse of the same frequency, it is said to be in resonance.

It is also possible for oscillation to result from periodical impulses that are not the natural frequency. This is generally called forced oscillation.

As you know, when the head/arm tracks the media surface, it is oscillating. If the media is extremely irregular, it is possible to meet the resonance of the head/arm assemblies. When this happens, control of these becomes difficult and they amplify the vibrations. Because the media is the closest contact, head crash results; if the media were not there, the system could destroy itself. An example of self-destruction is a crystal goblet, which, when the resonant frequency is reached, shatters.

Damping

Damping is a method used to help cut down on the resonance effect. An example of this is the shock absorber on a car. There is some natural damping built into the disk system at the time it's manufactured; however, if excessive damping was added, the response of the head/arm would be affected. Therefore, this damping is minimal, and, as noted before, preventive maintenance and care when handling the media are required.

Load/Unload Systems

Some fixed disk systems allow the heads to make contact with the media when the system is not in operation. This is accomplished by either moving the heads to the outer edges where there is no data, or, by applying a thin coating of a lubricant to the media. This lubricant is so thin that it does not impair data read/write, yet it prevents contact with the data track.

This activity deals only with the systems that have removable media and require the heads to be out of the way.

The read/write heads must be loaded to the disk surfaces before exchanging data with the controller. The heads must be removed (unloaded) from this position and driven clear of the disk pack when power is removed from the unit, when the disk pack velocity falls below a predetermined level, or when commanded to by the control logic. Remember, it is this velocity that keeps the air pressure high enough to allow the heads to fly. Two basic methods have been employed to load the heads, torsion arm and cam.

Torsion Arm

The torsion arm is not widely used today, in most cases it has been replaced by the cam method.

The head loading mechanism applies sufficient pressure to the read/write heads to hold them to within the 120 microinches of the disk surface (see figure 9-69).

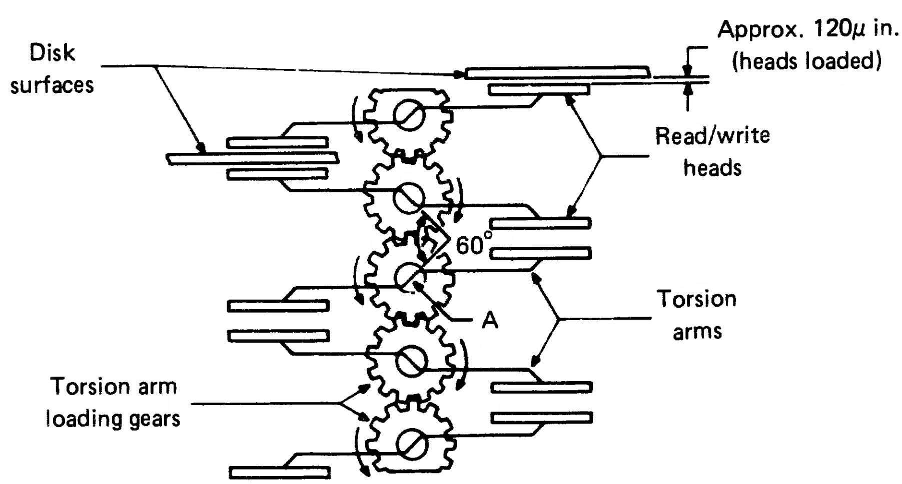

Figure 9-69. R/W heads and torsion arms with associated gears

This loading mechanism consists of the head loading cam, cam latch solenoid, cam release lug washer and spring, cam follower, center torsion and cam follower rod, the heads-loaded latch solenoid and switch, torsion arm gears, and the torsion spring (see figure 9-70).

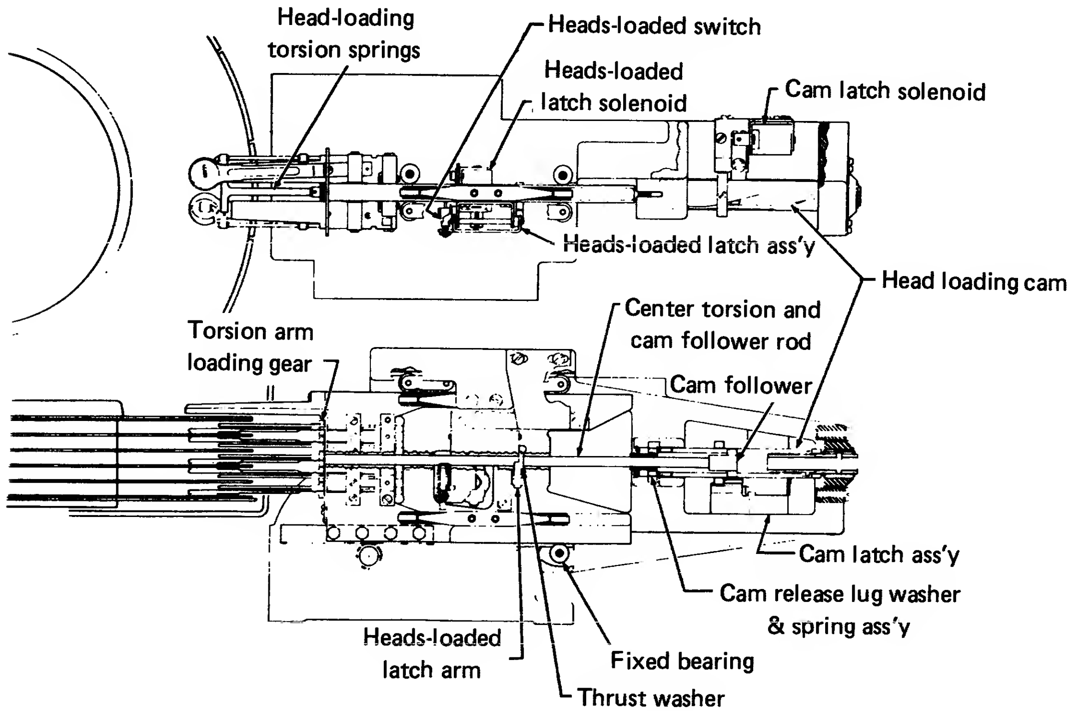

Figure 9-70. Head loading mechanism

When the power is off, the heads are retracted and unloaded to facilitate installation of a disk pack. When the power is turned on and the disk pack is UP TO SPEED, the heads are moved forward into the area between disk surfaces. Loading pressure is then applied.

Figure 9-71 shows a simplified head loading diagram.

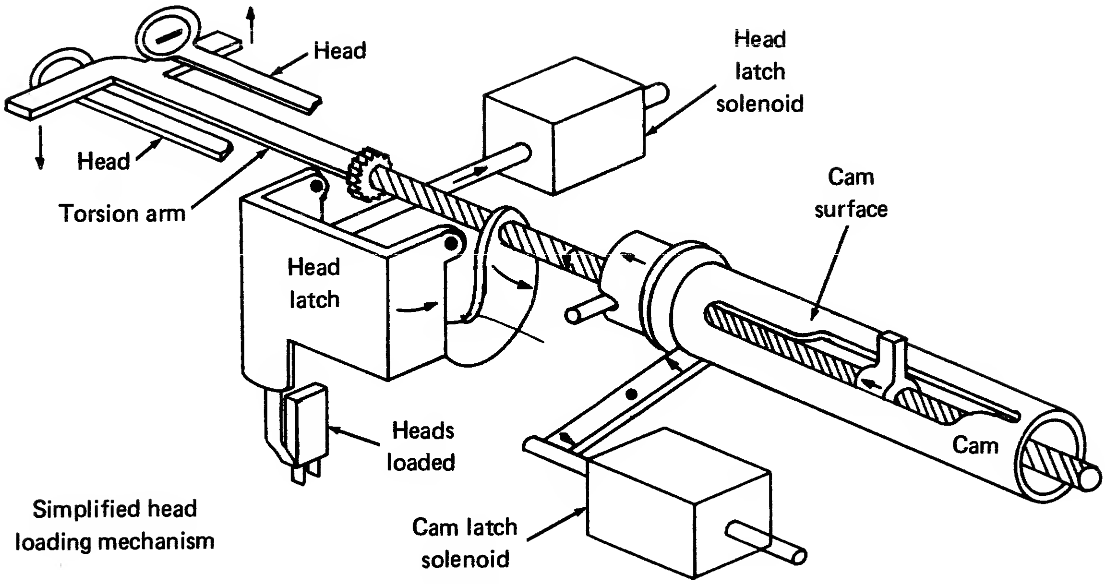

Figure 9-71. Simplified head loading mechanism

When power is sequenced up, the head loading cam is moved into a locked position by the cam latch solenoid. As the carriage moves forward, the cam follower rides the cam surfaces and rotates the center torsion and cam follower rod 60 degrees. This rotation is imparted through the middle torsion arm gear to the four remaining torsion gears.

A torsion arm, mounted to each gear, is connected between a pair of heads. By rotating the arm, one head is forced downward toward one disk surface and the other head is forced upward toward another disk surface.

At the end of the 60-degree rotation, the heads are loaded and the heads-loaded latch engages the center torsion and cam follower rod, locking it in place. The heads-loaded latch also transfers the HEADS-LOADED switch, which releases the cam latch solenoid.

The heads will remain loaded until power is lost or turned off. When power is turned off, the heads-loaded latch will release the center torsion and cam follower rod, allowing the heads to slowly unload under control of the air damper.

Cam Loading

Head loading amounts to allowing spring pressure of the floating arm (part of head/arm assembly) to move the aerodynamically shaped head face toward the related disk surface. When the cushion of air that exists on the surface of the spinning disk is encountered, it resists any further approach by the head. Spring pressure is designed to just equal the opposing cushion pressure (function of disk pack rpm) at the required height.

As a result, the head flies. However, if the spring pressure exceeds the cushion pressure (as would happen if the disk pack lost enough speed), the head stops flying and contacts the disk surface. This could cause damage to the head as well as the disk surface.

To prevent damage to the heads and/or the disk pack during automatic operation, loading occurs only after the disk pack is up to speed and the heads are over the disk surfaces. For the same reason, the heads unload automatically and are retracted if the disk pack rpm drops out of tolerance. During manual operations, heads should never be loaded on a disk pack that is not rotating.

The floating arm (figure 9-72) is designed to maintain a constant loading force. While the heads are retracted, head cams on the actuator housing bear against the floating arm cam surfaces. The cams support the loading force and hold the heads in the unloaded position. As the carriage moves forward, the floating arm cam surface rides off the head cam just after the read/write heads move out over the disk surface. The loading force moves the head face toward the air layer on the surface of the spinning disk until the opposing forces balance.

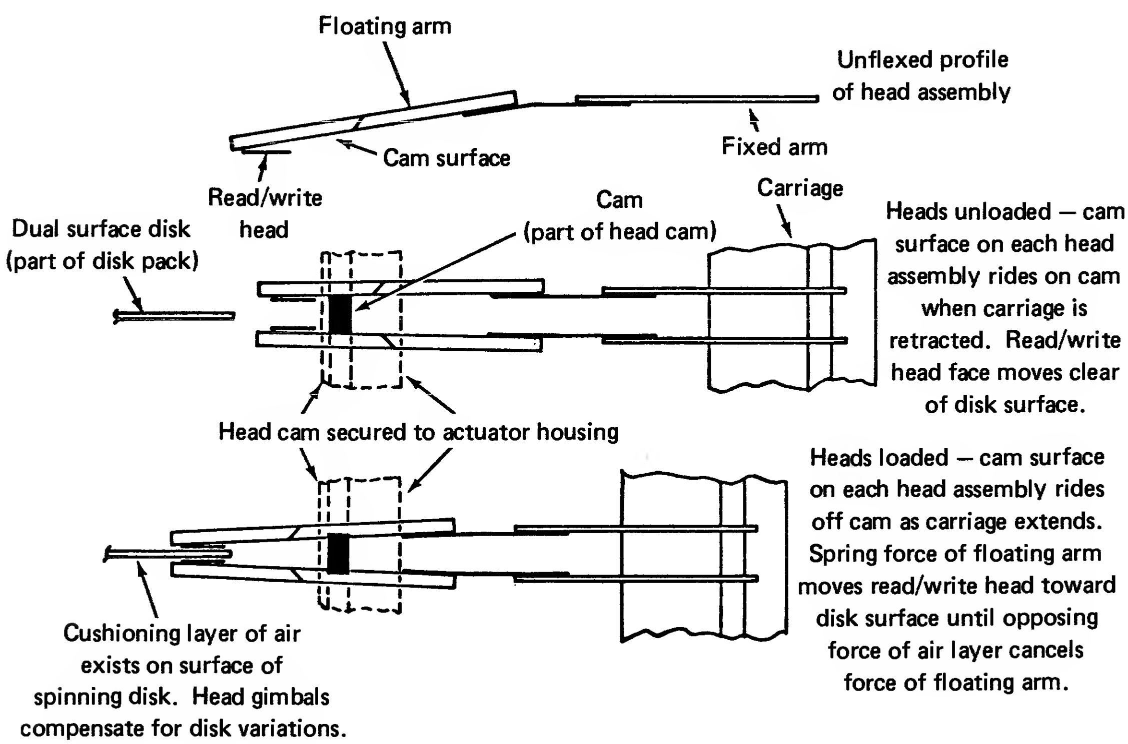

Figure 9-72. Head loading

The heads loaded switch status reflects the state of the read/write heads (loaded or unloaded). This status is used in the logic chassis and power supply. The switch mounts on the carriage track and is transferred by carriage motion. Whenever the carriage is fully retracted, the switch state reflects the unloaded status of the heads. As the carriage moves forward during a power on/first seek, the switch transfers at a point within about 0.1 to 0.2 inch forward of the retracted stop. This switch status remains unchanged until the carriage is retracted to the same position and, as such, does not precisely indicate the loaded/unloaded status of the heads.

Head unloading occurs whenever power to the unit is removed or disk pack rpm drops below tolerance. Either event drops a speed enable signal to the logic. This causes the voice coil to drive the carriage in reverse from its current location toward the retracted stop. As the carriage retracts, the cam surfaces encounter the head arms and each head rides vertically away from the related disk surface. The carriage continues back to the retracted position and stops.

Figure 9-73 shows the head/arm assembly motion.

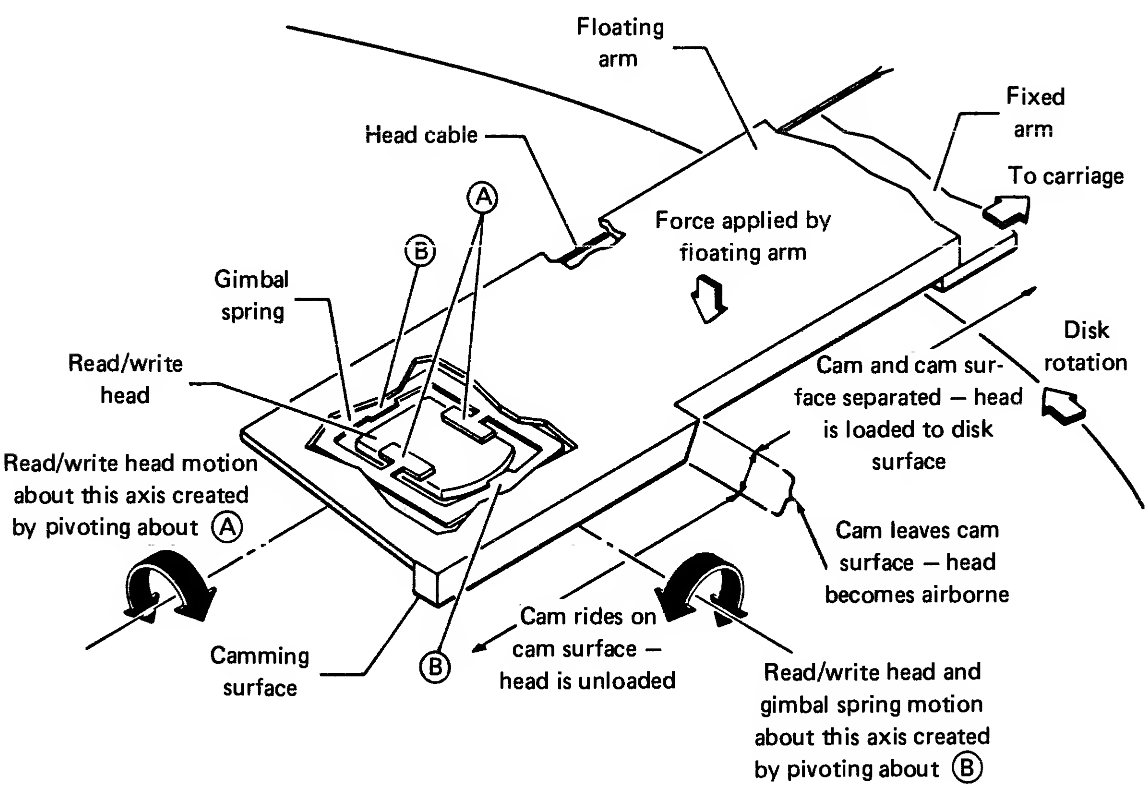

Figure 9-73. Head/arm assembly motion

Block 10

Movable Heads

Actuator Assemblies and Carriage Motion

Disk storage devices with movable heads need an actuator assembly to move the heads and position accurately. There are several different types of actuator assemblies; some that are still being used are no longer manufactured. Some actuators move the heads in a linear motion and others rotate the heads across the disk like a phonograph arm across a record. This activity gives functional descriptions of five different types of actuator assemblies and characterizes linear and rotary head motion.

Actuator Assemblies

Hydraulic Actuator

The hydraulic actuator used a hydraulic piston and cylinder assembly to move the carriage on which the heads were mounted. Disk storage devices with hydraulic actuators are no longer manufactured, although some may still be in use. See figure 10-1 for a simple functional diagram of a hydraulic actuator.

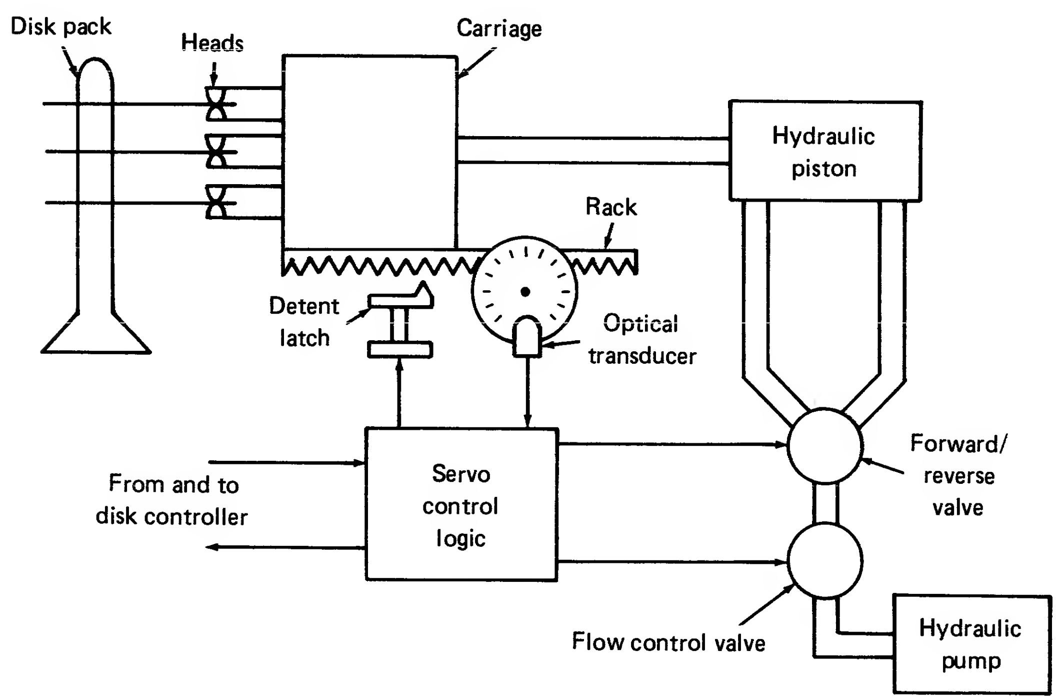

Figure 10-1. Hydraulic actuator

Head Movement

To move the heads in and out over the disks, the carriage is attached to a hydraulic piston. Oil flow from a hydraulic pump is regulated by flow volume control valves and directional valves to move the piston in and out. Servo control logic circuits interpret instructions from the disk controller and operate the flow control valves. A rack and pinion arrangement rotates a glass disk with etched lines; each etched line represents a track on the disk and an optical transducer reads the lines as they rotate by, sending that information to the servo circuits. When servo control logic senses that the heads are in the desired position, the detent latch is set into the teeth of the rack, locking the heads on track.

Fluid to Piston Control

Hydraulic actuators use two means of controlling the flow of fluid to the piston. One means, referred to as the "bang-bang" control, uses several flow valves, each with its own rate of flow. When the drive unit receives instructions to move the heads, all valves open at first, thereby accelerating the carriage. As the optical transducer counts tracks passing by and the heads come closer to the desired track, the valves close one at a time. This reduces the flow of fluid, and the speed of the carriage, until the heads are in position. Servo control logic then actuates the detent latch, locking the heads on track.

The other means of controlling fluid flow to the piston is proportional control, which uses one variable flow control valve. This method starts the heads moving with the valve wide open, and, as the heads approach the desired track, the flow control valve gradually shuts down the flow. Proportional hydraulic actuators are also used on position and velocity transducers. Servo control circuits use the position and velocity information to control the flow of fluid to the hydraulic piston.

Disadvantages

The hydraulic actuator became obsolete because its upper TPI (tracks per inch) limit was about 100 TPI. It also required a bulky, complex system of hydraulic pumps, pipes, and valves which added to the cost, size, and repair problems of the storage device.

Printed Circuit Motor Actuator

The printed circuit motor actuator used a high-speed rotary motor and a rack and pinion arrangement to move the carriage in and out. Printed circuit motor actuators are no longer produced. See figure 10-2 for a simple functional diagram of this type of actuator.

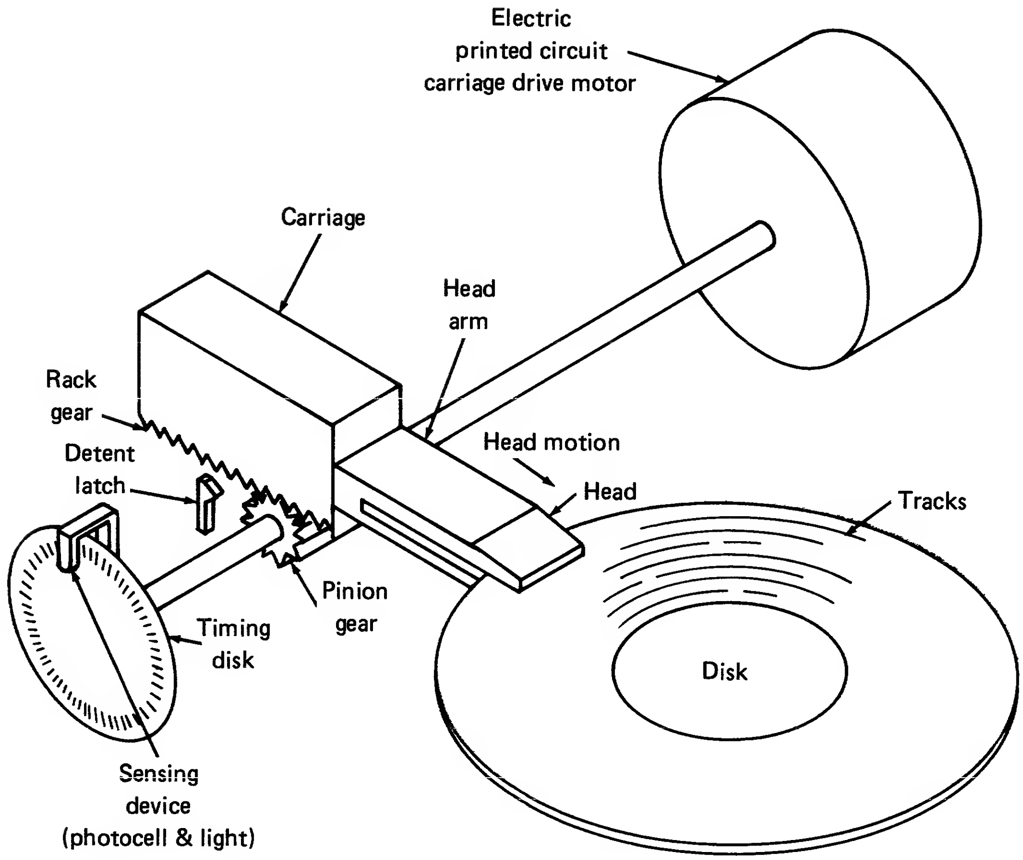

Figure 10-2. Printed circuit motor actuator

Head Movement

To move the heads in and out over the disks, the carriage is attached to a rack gear. A printed circuit motor rotates the pinion gear back and forth, moving the rack and carriage. A printed circuit motor is capable of quick acceleration and deceleration. Servo control logic circuits interpret instructions from the disk controller to regulate voltage going to the motor. The motor also rotates a glass disk with etched lines; it acts as a timing disk to indicate the position of tracks on the storage disks. An optical transducer senses the etched lines, then feeds that information back to the servo control circuits. When the circuits sense that the heads are over the desired track, the detent latch is set, locking the heads on track.

Disadvantage

Like the hydraulic actuator, the printed circuit motor actuator was limited to about 100 TPI.

Linear Voice Coil Actuator

The linear voice coil actuator utilizes the principle that by applying current to a coil within a magnetic field, the coil moves in or out of the field, depending on the polarity of the current applied to the coil. The voice coil actuator represented a major advance over the hydraulic and printed circuit motor actuators and is the mechanism used in virtually all disk storage devices with movable heads manufactured today. See figure 10-3 for a simple functional diagram of a linear voice coil actuator.

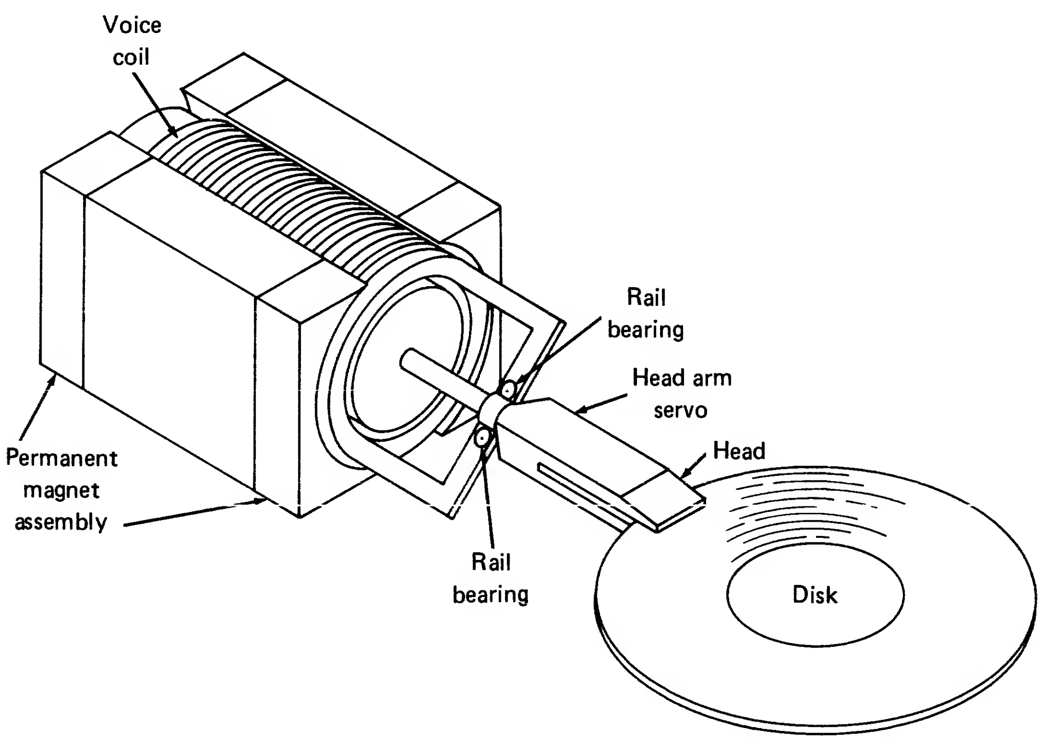

Figure 10-3. Linear voice coil actuator

Head Movement

To move the heads in and out over the disks, the carriage is attached to a coil of wire called the voice coil. (It is called a voice coil because it is similar to the coils in the backs of permanent magnet speakers.) The voice coil moves back and forth within the air gap of a powerful permanent magnet; the magnet generates a strong magnetic flux across the air gap. Electric current moving through the coil interacts with the magnetic field across

the gap, generating a force on the coil which is capable of moving the carriage and the heads. Direction and velocity of head movement is directly proportional to the polarity and amount of current moving through the coil.

Current is supplied to the voice coil by a power amplifier which amplifies low voltage signals from the servo control circuits. The servo control circuits rely on instructions from the disk controller and feedback information on head position and head velocity to generate control signals going to the voice coil. The voice coil also acts to hold the heads on the correct track.

Advantages

The voice voil actuator has several advantageous characteristics. It has only one moving part, the coil, compared to the many moving parts in the hydraulic and printed circuit motor actuators. There are no touching parts and no friction in the actuator itself, although the carriage rides on rails and bearings. These two features make the voice coil actuator quicker and more reliable than previous devices. The voice coil also has very quick response and is much more accurate than other actuators. TPI figures now reach to almost 500 TPI and may go further because of the accuracy of the linear voice coil.

Rotary Motion Voice Coil

The rotary motion voice coil actuator uses the same electromagnetic principle as the linear voice coil except that the heads are moved in an arc across the data tracks in the same manner that a phonograph arm moves across the tracks of a record. Rotary motion voice coil actuators are often used on fixed pack disk storage devices. Once thought of as a replacement for linear voice coils, rotary motion voice coil actuators are now taking their place in specific applications in the disk storage field. See figure 10-4 for a simple functional diagram of a rotary motion voice coil actuator.

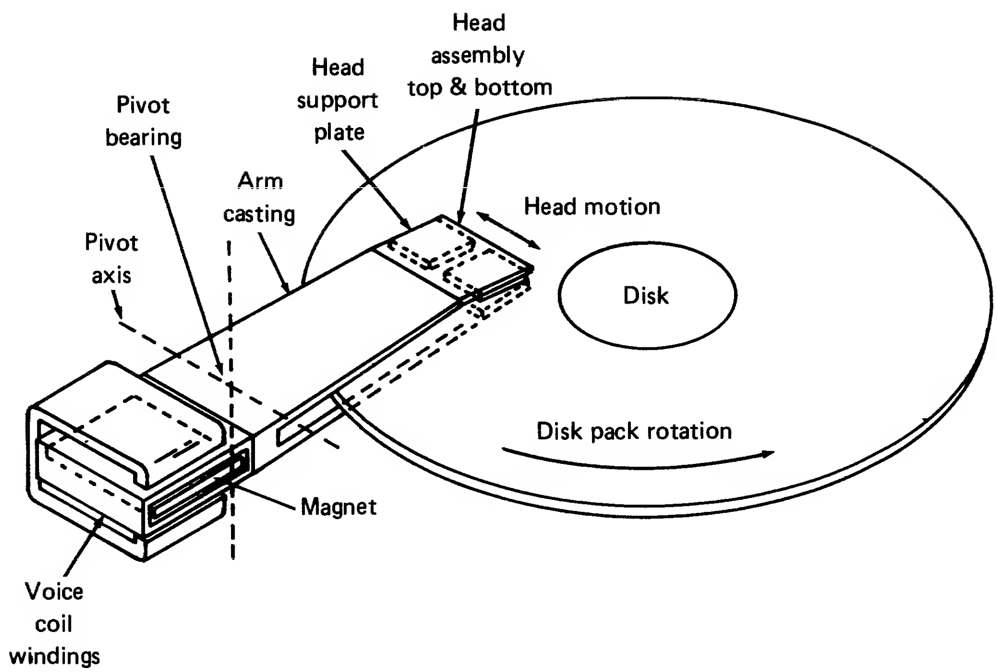

Figure 10-4. Rotary motion voice coil actuator

Head Movement

The heads are mounted on a carriage which pivots back and forth across the data tracks like a phonograph arm. However, the data tracks are still circular, not spiral like phonograph record tracks. A rectangular voice coil is mounted on the end of the carriage opposite the heads; the voice coil fits over a strong permanent magnet and supplies force to the end of the carriage in the same way that a linear voice coil operates. The rotary voice coil has circuitry similar to a linear voice coil and is controlled in the same manner.

Advantages/Disadvantages

The primary advantage of the rotary voice coil is its lower cost than linear voice coil mechanisms. The bearing arrangement supporting the carriage/arm is simpler than the bearing/rail arrangement supporting the carriage of a linear voice coil. However, this same bearing arrangement makes the rotary voice coil slightly less accurate than a linear voice coil, reducing the number of tracks per inch that can be recorded.

Stepper Motor Actuator

The stepper motor actuator, at first, seems to be a step backward in actuator technology because it uses a motor and gear arrangement to move the carriage and heads. It is relatively slow and inaccurate. However, when used with a floppy disk, it is a very appropriate technology for low-cost, low-capacity disk storage. See figure 10-5 for a simple functional diagram of a stepper motor actuator.

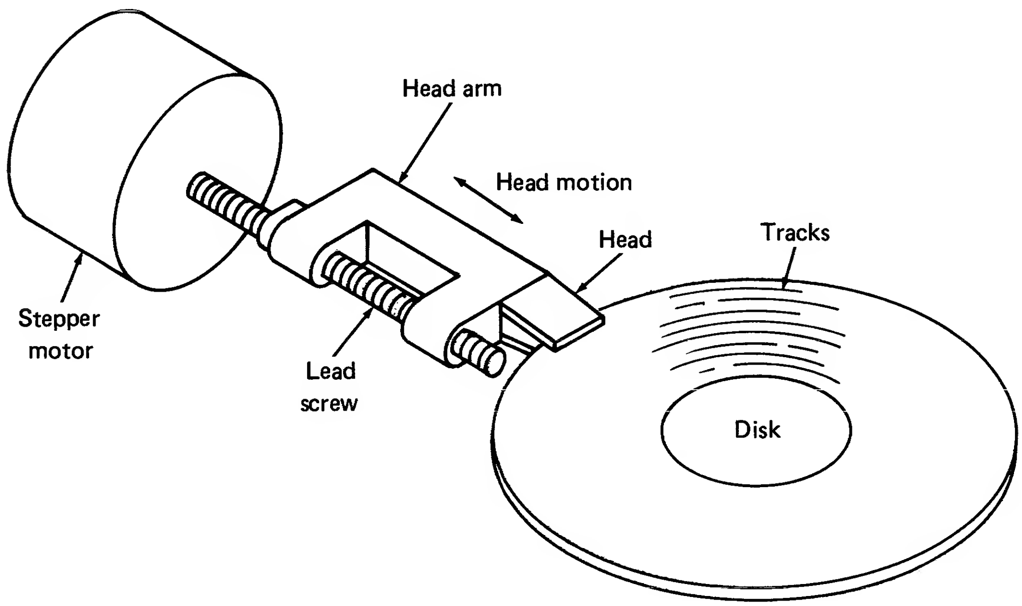

Figure 10-5. Stepper motor actuator

Head Movement

The stepper motor itself is the key element of this type of actuator. When the motor receives a control signal or pulse, it turns only a specific number of degrees, and it will turn the same number of degrees for each pulse it receives. The motor turns a lead screw, which in turn moves the head in and out; for each pulse, the head will move only a specific distance, usually one track. In this way, the servo control logic for the stepper actuator moves the heads a specified number of tracks simply by sending the same number of pulses to the stepper motor.

The lead screw arrangement produces linear motion; other types of stepper motor actuators use a rack and pinion arrangement to produce linear motion. A later model stepper motor actuator produces rotary motion by wrapping a taut metal band around a wheel attached to the stepper motor. As you can see in figure 10-6, when the wheel pulls the band back and forth, the head is pulled back and forth in a rotary motion.

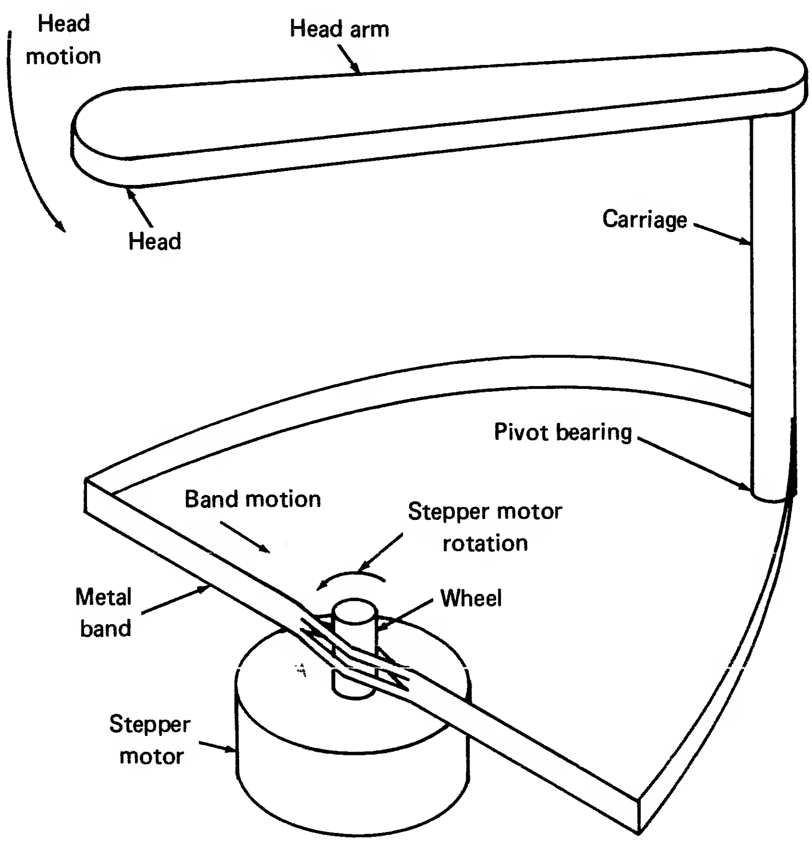

Figure 10-6. Rotary motion produced by band mechanism

Advantages/Disadvantages

The primary advantage of the stepper motor actuator is its low cost. Its mechanism is relatively simple and the logic circuits needed to control it are far simpler than those required to control voice coil actuators. However, its slow access time and poor accuracy limit its use to floppy disks where cost, not speed and capacity, is the important factor to the customer.

Carriage Motion

The actuators described above illustrate two types of carriage motion: linear and rotary. Linear motion actuators, such as the hydraulic, voice coil, and some stepper motor actuators, move the heads straight in and out of the disk pack. The rotary motion voice coil and some stepper motor actuators swing the heads in an arc across the disk tracks.

Each type of motion has its own characteristics and advantages. Actuators relying on linear motion can generally use sturdier carriages on which many heads can be mounted. Disk storage devices using four, five, or more disks per pack generally use linear actuators because a rotary actuator for that many disks would tend to vibrate at such a low natural frequency that it would impair the positioning accuracy. With linear actuators, the heads always maintain the same angle with respect to the track. However, with rotary actuators, the heads are always changing their angle with respect to the data tracks, complicating read and write methods.

Rotary motion actuators are more compact than linear motion actuators and require less space outside the disk pack. Linear guide rail motion actuators require long, extremely straight bearings for the carriage to ride on while rotary motion carriages ride on a simple set of bearings placed between the heads and the voice coil.

Summary

There is a wide variety of actuator mechanisms that move the heads on disk storage devices, although the types most widely used now are the linear motion voice coil and rotary motion voice coil. Stepper motor actuators are common in low-cost, low-capacity floppy disk devices.

The actuators rely on two types of carriage motion, linear and rotary. Each type of motion has its own advantages and disadvantages.

Velocity Transducer

A seek operation on a modern disk storage device is a complex procedure; there is more to it than simply moving the heads from one cylinder to another cylinder. Efficient operation requires that the device perform seek operations as quickly as possible to reduce access time. But the velocity of the heads and carriage must be closely monitored and regulated to prevent the inertia of the carriage from carrying the heads beyond the desired cylinder. To accomplish this, all disk storage devices, except those with stepper motor actuators, need some kind of velocity transducer. This activity describes the basic concepts behind velocity transducers and identifies two types of velocity transducers.

Basic Concepts

Modern disk storage devices rely on proportional control to reduce access time to a minimum. Proportional control means that for a seek of any number of cylinders, there is an ideal velocity curve that involves rapid acceleration of the carriage, then smooth deceleration, stopping the heads on the desired cylinder. Figure 10-7 illustrates such a velocity curve. The curve is designed by considering the mass of the carriage and capabilities of the actuator and its associated electronic circuitry. At any point on the curve, the heads are traveling at optimum speed.

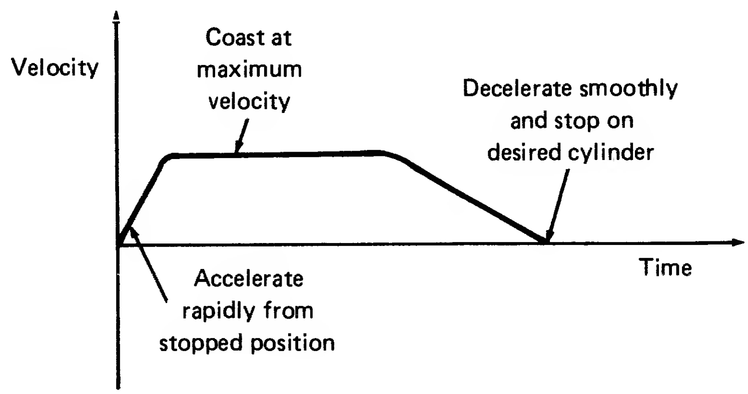

Figure 10-7. Velocity curve

The desired velocity curve is generated by servo control circuits in the disk storage device. During a seek operation, the velocity curve is represented by a voltage level that indicates what the velocity of the carriage should be during each portion of the seek. The velocity transducer generates a voltage level that is proportional to the velocity of

the carriage. By comparing the desired velocity voltage and the actual velocity voltage, the servo control circuits can tell if the carriage needs to be accelerated or decelerated. If the desired velocity voltage is greater than the actual velocity voltage, current is supplied to accelerate the carriage. If the actual velocity voltage is greater than the desired velocity voltage, meaning that the carriage is going too fast, braking current is applied to the actuator.

Disk storage devices that use the servo disk/servo head position transducer have an additional use for the velocity transducer. Not only is it used for velocity control during the seek, but it is also used to dampen movement of the heads once they are on cylinder and following the track.

As mentioned above, velocity transducers have one simple, basic characteristic: the voltage they produce must be proportional to the velocity of the carriage. This means that the signal level must increase as velocity increases and must change polarity when the direction of the motion changes. There are two types of velocity transducers: the linear velocity transducer and the electronic tachometer.

Linear Velocity Transducers

Linear velocity transducers rely on electromagnetic principles to measure the velocity of the carriage. They consist of a thin, pencil-shaped permanent magnet which moves freely within a long coil of wire. Figure 10-8 illustrates this relationship.

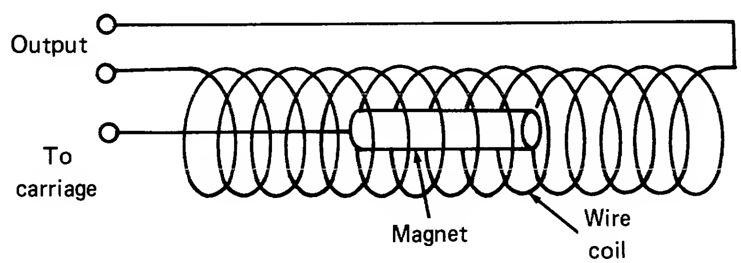

Figure 10-8. Linear velocity transducer concept

As the magnet moves within the coil, its lines of flux induce voltage in the coil. This voltage is proportional to the velocity of the magnet and it changes polarity when the magnet changes direction.

The magnet is linked to the carriage and travels with it. On linear voice coil actuators, the velocity transducer is sometimes mounted within the voice coil itself. With other types of actuators, the transducer is mounted elsewhere on the deck.

Electronic Tachometers

Electronic tachometers do not use a separate, mechanical transducer to actually measure the velocity of the carriage. Rather, they have circuitry that inputs 1) the current level being supplied to the actuator, and 2) the position signal. The circuits interpret these two inputs and produce a voltage that is proportional to the velocity of the carriage. The output of the electronic tachometer is functionally identical to the output of the linear velocity transducer. Figure 10-9 illustrates the inputs and outputs of the electronic tachometer. Although the electronic tachometer could conceivably work with any actuator that has a position transducer, it is almost always used with disk storage devices that have a voice coil actuator and servo disk/servo head position transducer.

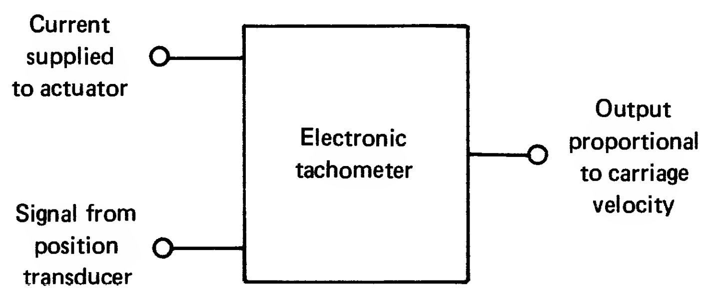

Figure 10-9. I/O electronic tachometer

The electronic tachometer works by summing the current supplied to the actuator and differentiating the position signal. Figure 10-10 shows that the velocity of the carriage is the integral (the result of a mathematical integration) of the current supplied to the actuator.

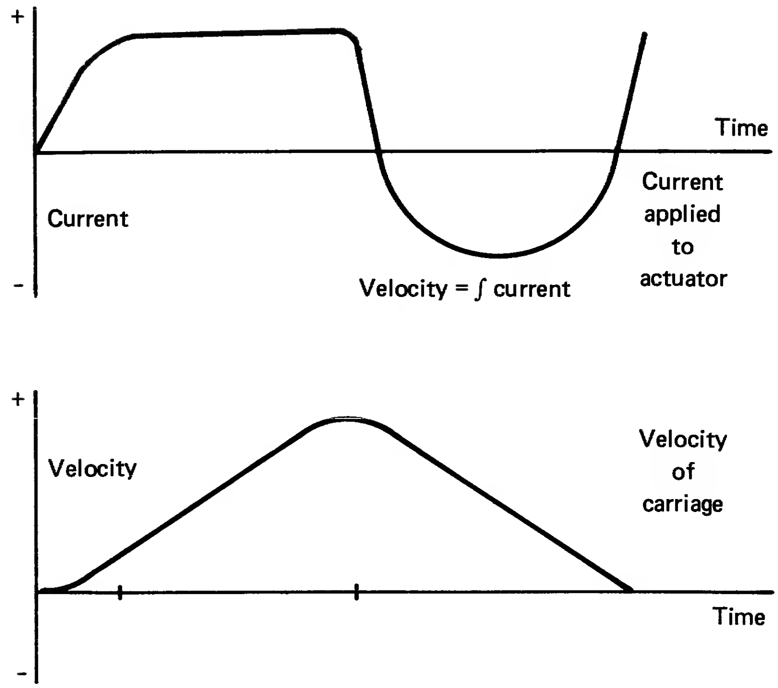

Figure 10-10. Current/velocity relationship

However, some small part of the current applied to the actuator is spent in overcoming friction and windage, but does not produce acceleration. This introduces a small inaccuracy. To correct this inaccuracy, the slope of the position signal is used to give a velocity signal near each cylinder crossing that is valid.

Figure 10-11 shows the position signal. The slope of the signal around the null, when differentiated, gives an accurate indication of the velocity of the carriage.

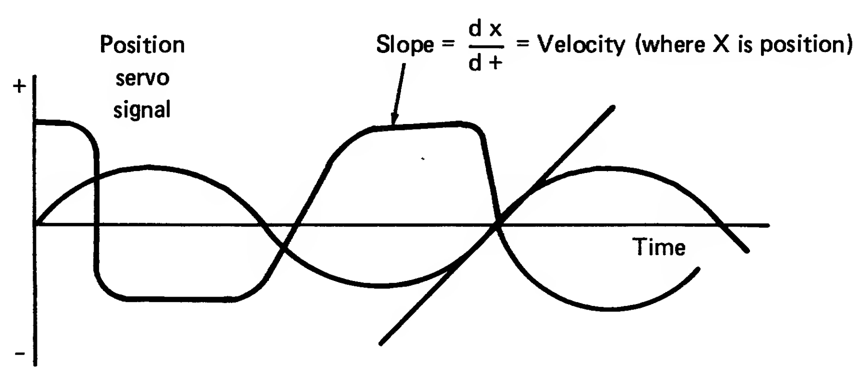

Figure 10-11. Position signal

Summary

Velocity transducers are needed to keep carriage velocity at an optimum level during seek operations and to reduce access time to a minimum. Velocity transducers produce a voltage that is proportional to the velocity of the carriage and that changes polarity when carriage motion changes direction. There are two common types of velocity transducers: linear velocity transducers and electronic tachometers. Linear velocity transducers measure carriage velocity with a magnet moving in a coil of wire. Electronic tachometers derive the carriage velocity by interpreting current supplied to the actuator and position signals.

Position Transducer

In the process of performing a seek operation, the actuator starts by moving the heads from their present cylinder toward the desired cylinder. However, to position the heads exactly on the desired cylinder, the servo control logic circuits need constant information regarding the position of the heads. This information is supplied by the position transducer. This activity describes the basic concept behind position transducers and gives several examples of position transducers.

Basic Concepts

A position transducer does not directly indicate the position of the heads on the disk pack. Rather, it generates a signal that indicates the position of the heads relative to one single cylinder. Generally, this position signal is bipolar, so that when the heads are on one side of the cylinder, the signal is, for example, negative. As the heads cross the cylinder, the signal goes to zero and then to positive when heads are on the other side of the cylinder. This type of signal is illustrated in figure 10-12.

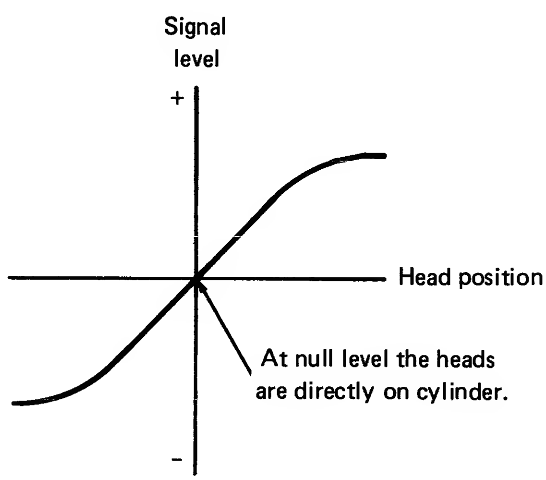

Figure 10-12. Cylinder position signal

This bipolar signal can be used two ways: to count cylinder crossings or to keep the heads on cylinder. Figure 10-13 shows the signal that is generated as the actuator moves the heads across the disks.

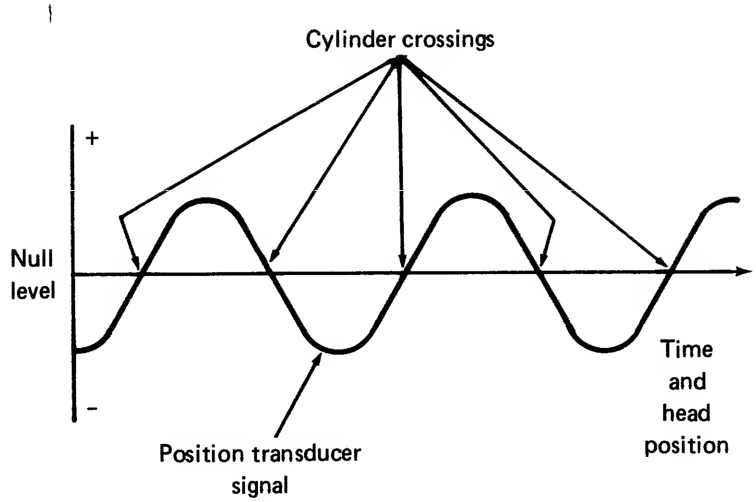

Figure 10-13. Position signals across disk

Each point where the signal crosses the null level indicates the crossing of a cylinder. Note that it doesn't indicate which cylinder is crossed. But the servo control logic circuits can sense the null level and count the number of times it occurs.

At the beginning of a seek operation, a disk storage device usually has the present cylinder number and the desired cylinder number stored in its memory registers. From these two numbers, the device can determine how many cylinders must be crossed to reach the desired cylinder. The actuator starts moving the heads, and the position transducer starts generating the waveform (see figure 10-13). Servo control circuits count the nulls. When the count equals the number of the cylinder that the head must be moved to, the actuator stops and the heads are locked in position.

Look at figure 10-12 again. If the heads are on cylinder, the position transducer generates a NULL signal; but if the heads drift off cylinder, the transducer generates either a POSITIVE or NEGATIVE signal. Servo control circuits can sense this ERROR signal and respond by making the actuator move the heads back toward the cylinder. When the ERROR signal is back to zero, the actuator stops and the heads are back on cylinder.

Essentially, the position transducer generates a signal that varies from one polarity to another and is null when the heads are directly on a cylinder. This signal can be used to count cylinder pulses or to keep the heads on cylinder. However, there are several different types of position transducers that generate that signal.

Optical Transducers

Optical transducers generally use a piece of glass with a grid pattern etched on it, a light source, and a light sensor with a mask that matches the grid pattern on the piece of glass. The piece of glass moves with the carriage, passing between the light source and sensor. The etched lines on the glass periodically interrupt the light, making the photoelectric sensor generate a triangular or sinusoidal signal that corresponds to the tracks on the disk. There are two types of optical transducers: the linear optical transducer and the rotary optical transducer (see figures 10-14 and 10-15). However, this does not mean that linear transducers are used only on linear actuators and rotary transducers only on rotary actuators.

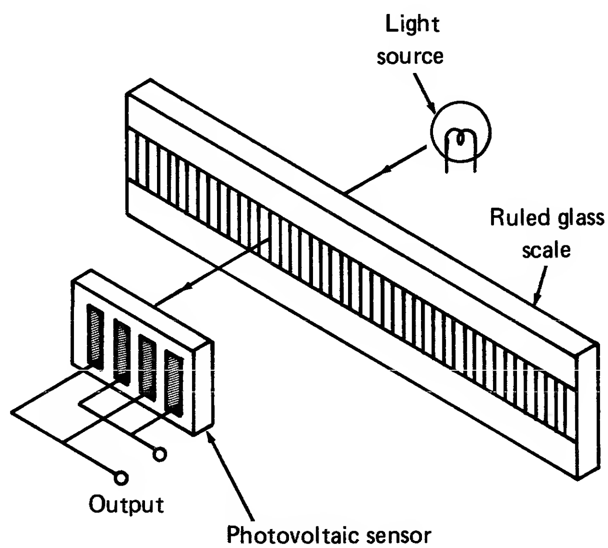

Figure 10-14. Linear optical transducer

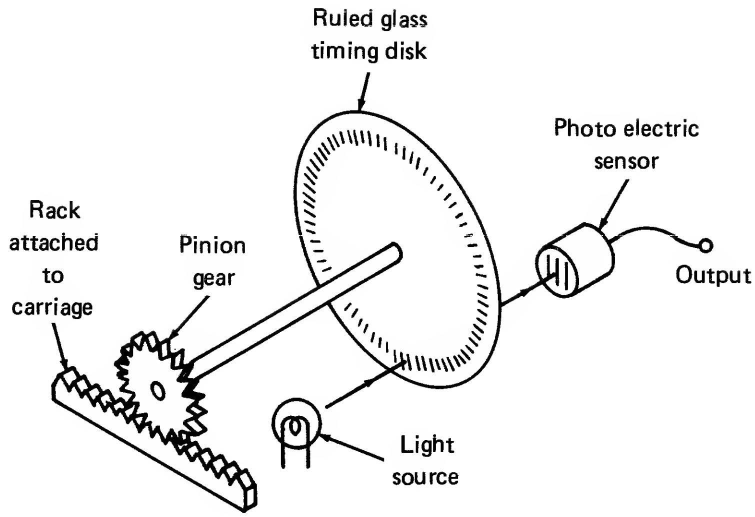

Figure 10-15. Rotary optical transducer

The linear transducer moves directly with the carriage, but the rotary transducer needs a rack and pinion arrangement to translate the linear movement of the carriage into rotary motion for the timing disk.

The mask on the moving piece of glass alternately allows a maximum amount of light to fall on the photocell and then blocks it entirely. By making the midrange point between the two extremes the null point, the photocell generates the bipolar signal needed by the servo control circuits. The transducer is adjusted so that the null occurs when the heads are directly on cylinder.

Advantages/Disadvantages

Optical transducers are low-cost, reliable devices for indicating the position of the heads, and their associated electronic circuits are simple and inexpensive.

Optical transducers have two disadvantages. First, glass has a different thermal coefficient of expansion than aluminum, the material used to make the recording disks. This means that as the operating temperature of the storage device changes, the disk and transducer will expand at different rates, throwing them out of proper adjustment. This limits the accuracy of the head-to-track positioning.

Second, since the optical transducer is an external transducer (meaning it is not mounted on the disk pack itself), the markings on the transducer may not conform exactly with the cylinders on the disk pack. This is a particular problem when disk packs are moved from machine to machine, with each machine adjusted somewhat differently.

Even with these limitations, optical transducers are still used on some low-cost disk storage devices.

Inductive Transducers

Inductive transducers are frequently called "inductosyns." Figure 10-16 illustrates the two parts of an inductive transducer, the scale and the slider.

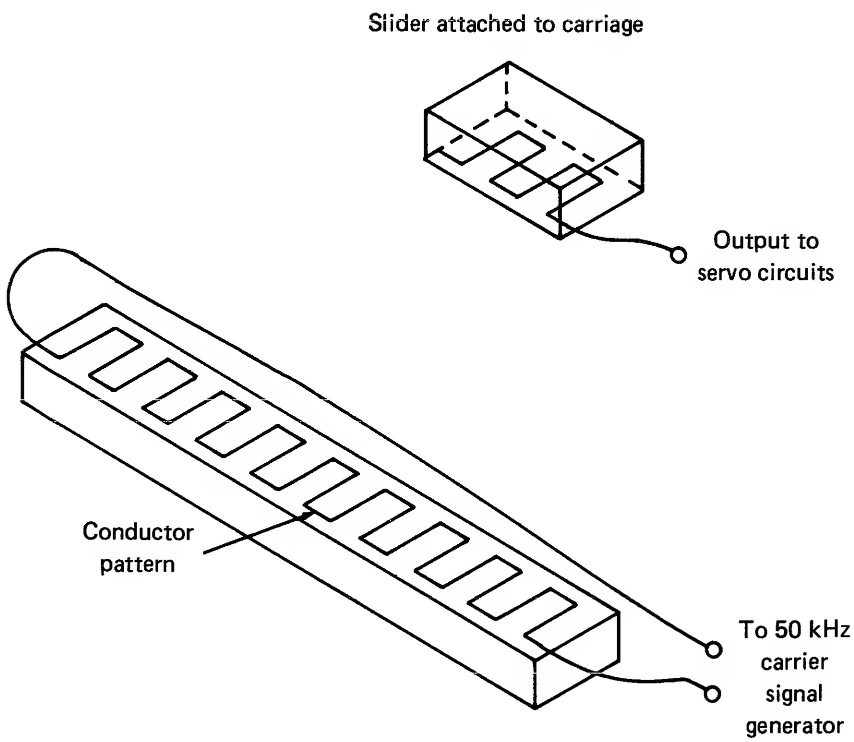

Figure 10-16. Inductive transducer scale and slider

Both the scale and slider have a printed circuit grid pattern etched onto their surfaces. The grid of the scale is excited with a 50 kHz carrier signal which has a potential to induce a signal in the grid of the slider.

The slider moves with the carriage, sliding over the grid surface of the scale. When the grid pattern on the slider matches up with the grid pattern on the scale, a maximum signal is induced in the slider. When the lines on the slider lie between the lines of the

scale, a minimum signal is induced in the slider. As the slider moves along the scale, it generates the sinusoidal signal needed by servo control circuits to count cylinder crossings. The scale is adjusted so that the nulls will occur when heads cross cylinders on the disk pack.

Advantages/Disadvantages

Inductosyns have an advantage over optical transducers in that they can be bonded to aluminum and have the same coefficient of expansion as the disk pack, making them more accurate.

However, the device, along with its associated circuitry, is more complex and expensive. Like optical transducers, inductosyns are external transducers with limited accuracy, creating problems when disks are transferred from machine to machine. Neither optical transducers nor inductive transducers can be used to keep the heads on track.

Servo Disk and Servo Head

In the servo disk/servo head system, the position transducer scale is recorded as magnetic tracks on an extra disk surface in the disk pack. The transducer pickup is an extra head mounted on the carriage with the other heads. Since the scale is right on the disk pack, and the sensor moves along with the heads, the problem of different coefficients of expansion is eliminated. Also, the scale moves with the disk pack, so compatibility between machines is enhanced. Given these features, this system is much more accurate than other position transducers and has enabled the design of disk storage devices with very high TPI capabilities. Figure 10-17 is a simple, functional illustration of the servo disk/servo head system.

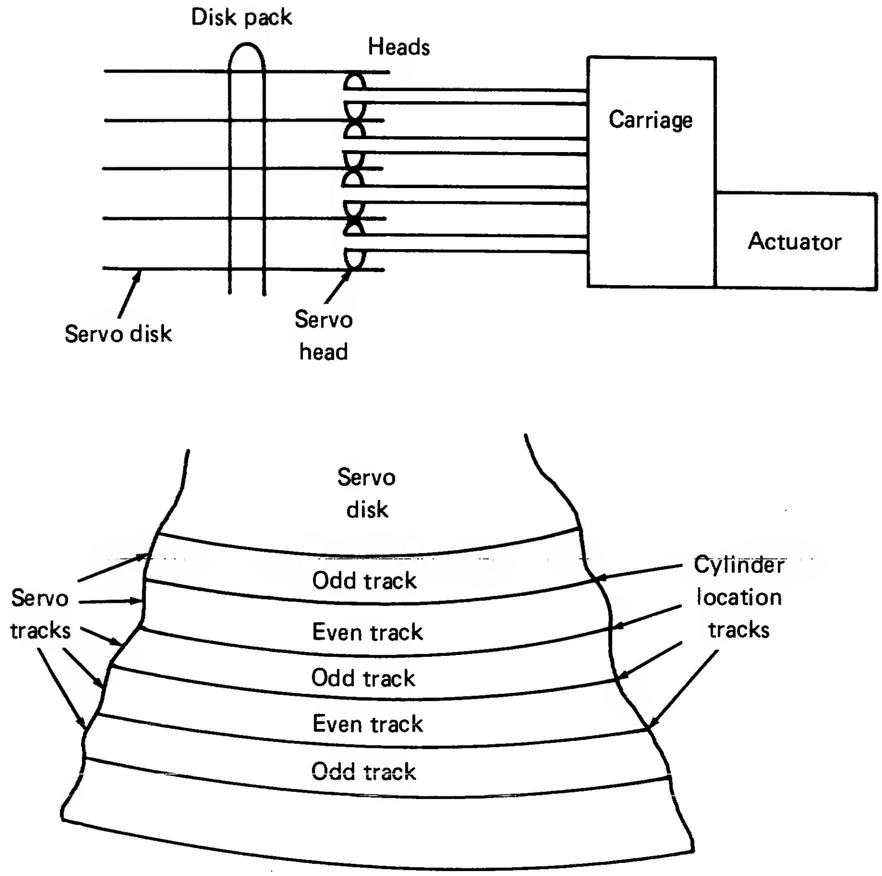

Figure 10-17. Servo disk/servo head system

The servo disk has tracks of magnetic pulses recorded on it at the time of manufacture. There are two different types of tracks, called odd tracks and even tracks, recorded alternately. The edge of each track is very close to the edge of the other, with only a very small gap between them. The edge, where two different types of tracks meet, marks the cylinder where data tracks will be recorded on all other disks in the pack.

On disk storage devices using the servo disk/servo head system, there are several different types of signals recorded on the odd and even servo tracks. Depending on the pattern of pulses, they are called Dibit, Tribit, or Quadbit tracks. In general, however, the servo tracks are recorded so that, when read by the servo head, the odd and even tracks generate signals that are the inverse of the other.

The servo control circuits interpret the signals from the servo head to produce a signal that is of one polarity when the head is over only an odd track; the signal is of opposite polarity when the head is over an even track, and is null when the head is directly over

the edge between odd and even tracks. This is the null needed to indicate when the heads are on cylinder. Figure 10-18 shows the servo signal generated as the servo head crosses the servo disk.

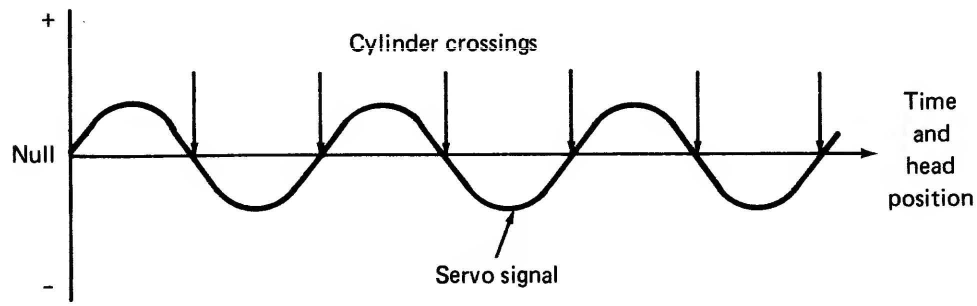

Figure 10-18. Servo signals across the disk

As with the other position transducers, the nulls can be counted for cylinder crossings. But the servo disk/servo head positioning system is extremely accurate. It can keep the heads on cylinder with very little error. Whenever the servo circuits sense the slightest variation from null, they can instruct the actuator to correct the error. The disk controller can even instruct a disk storage unit with a servo disk to move the heads slightly off track, if necessary, to recover marginal data.

Advantages/Disadvantages

Once again, the advantages of the servo disk/servo head system are its accuracy and excellent compensation for thermal expansion. However, its disadvantages are in terms of cost. The system requires an extra head and extra disk surface. The circuits that are required for interpreting the odd and even track signals into a bipolar error signal are complex and expensive.

Despite these disadvantages, the servo disk/servo head system is very common and almost universal on high-cost, high-capacity disk storage devices.

Summary

Position transducers are needed to position the heads in the disk pack. In general, they work by generating a bipolar signal that is null when the heads are on cylinder. The bipolar signal can be generated by several types of transducers but the servo disk/servo head transducer system is the most accurate.

Closed and Open Loop Servo

Disk storage devices are essentially "slave units"—they take instructions from the CPU through the disk controller. Any unit that carries out instructions from another unit is called a servo. On disk storage devices, the actuator and its associated electronic control circuits is a servo because it carries out the instructions to find a cylinder address on the disk pack.

There are two types of servos: closed loop servos and open loop servos. Open loop servos carry out instructions without checking to determine if they have performed the operation correctly. Closed loop servos constantly check their performance to make sure that the instructions are carried out correctly.

This activity lists the characteristics of both open and closed loop servos and gives functional descriptions of how each works.

Open Loop Servo

Among the actuator assemblies described in an earlier activity, the only actuator that is an open loop servo is the stepper motor actuator. Remember, a stepper motor is a special type of motor that turns only a specific number of degrees each time it receives a pulse. In a stepper motor actuator, the motor drives a lead screw, rack and pinion, or wheel and metal band to move the heads. The control logic performs a seek operation by applying a specific number of pulses to the stepper motor. Figure 10-19 is a simple block diagram of the stepper motor actuator. It illustrates the characteristics and functions of an open loop servo.

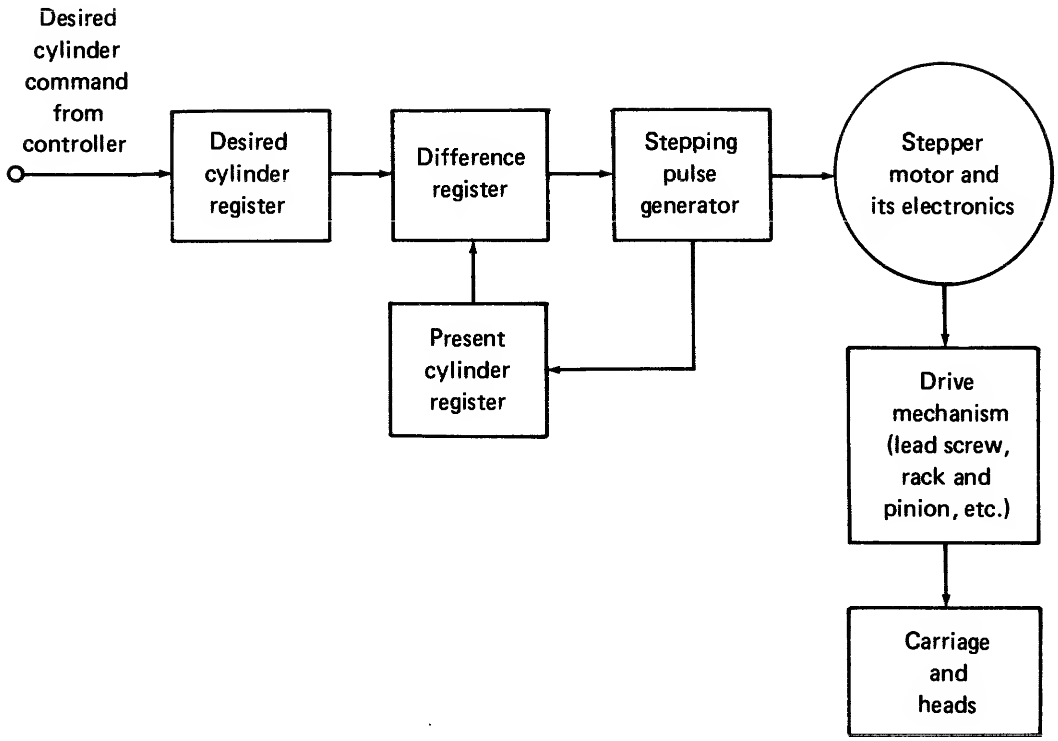

Figure 10-19. Stepper motor actuator

The disk controller commands the device to move the heads to the new, desired cylinder; that cylinder number is placed in a register. The difference register compares the desired cylinder with the cylinder the heads are presently on. If there is a difference, the difference register signals the stepping pulse generator to send a pulse to the stepper motor. The motor steps, in the direction also determined by the difference register, and moves the heads one track. At the same time, the stepping pulse generator also updates the present cylinder register so that it continues to hold the cylinder that the heads are presently on.

Again, the difference register compares the desired cylinder with the present cylinder, and, if there is a difference, instructs the stepping pulse generator to pulse the stepper motor and move the heads one track. This continues until the present cylinder is the same as the desired cylinder and the seek operation is complete.

The key characteristic of this open loop servo is that there is no indication of what the heads are actually doing. As an extreme example, the heads could have fallen off at the beginning of the seek, but the stepper would continue to drive the mechanism. As a

more realistic example, the drive mechanism could be slightly out of adjustment. As a result, the heads would be slightly off track—but there is no way the servo would have of sensing or correcting this error. A secondary characteristic of the open loop servo is that it has the potential to be very inaccurate. However, it is simple and relatively inexpensive.

Figure 10-20 illustrates an open loop servo a little more abstractly. The servo receives an instruction and the servo electronics interpret the instruction and signal the actuator to carry out the instruction in a discrete number of steps. There is no feedback to compare the actual results of the operation with the original instructions.

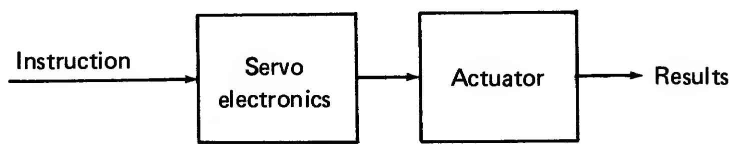

Figure 10-20. Open loop servo

If you object and point out that there is a loop in the stepper motor servo, you would be correct. In figure 10-19 you see there is a loop from the difference register, to the stepping pulse generator, to the present cylinder register, and back to the difference register. However, this loop still does not provide any information on what the heads are actually doing.

Closed Loop Servo

Figure 10-21 shows how the open loop servo in figure 10-20 can be changed to a closed loop servo by simply adding a feedback line.

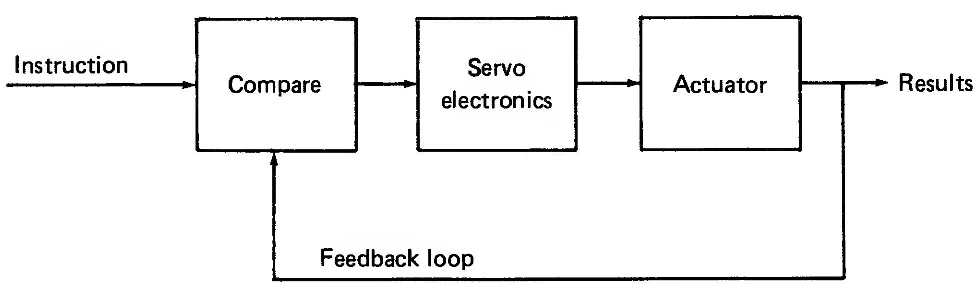

Figure 10-21. Closed loop servo basic block diagram

Now the instructions to the servo can be compared with the actual results of the operation and the actuator can continue to operate until the results match the instructions, within some reasonable tolerance. This illustrates the key characteristic of the closed loop servo: the results of an operation can be compared with the original instructions, and the servo can continue to operate until the results come within some specified tolerance of the instructions. As a result, the closed loop servo is much more accurate than the open loop servo, although it is more complex and expensive.

The key thing to remember in the operation of the closed loop servo is that it does not carry out instructions in discrete steps (like the open loop servo). It operates continuously until the results match the instructions within the specified tolerance. This means that the open and closed loop servos not only function differently, but also malfunction differently. An open loop servo may simply carry out an instruction incorrectly, but a closed loop servo may be unable to carry out an instruction within the specified tolerance and begin oscillating back and forth, trying to carry out the instruction.

Figure 10-22 expands on figure 10-21 to show how the closed loop servo functions on a disk storage device.

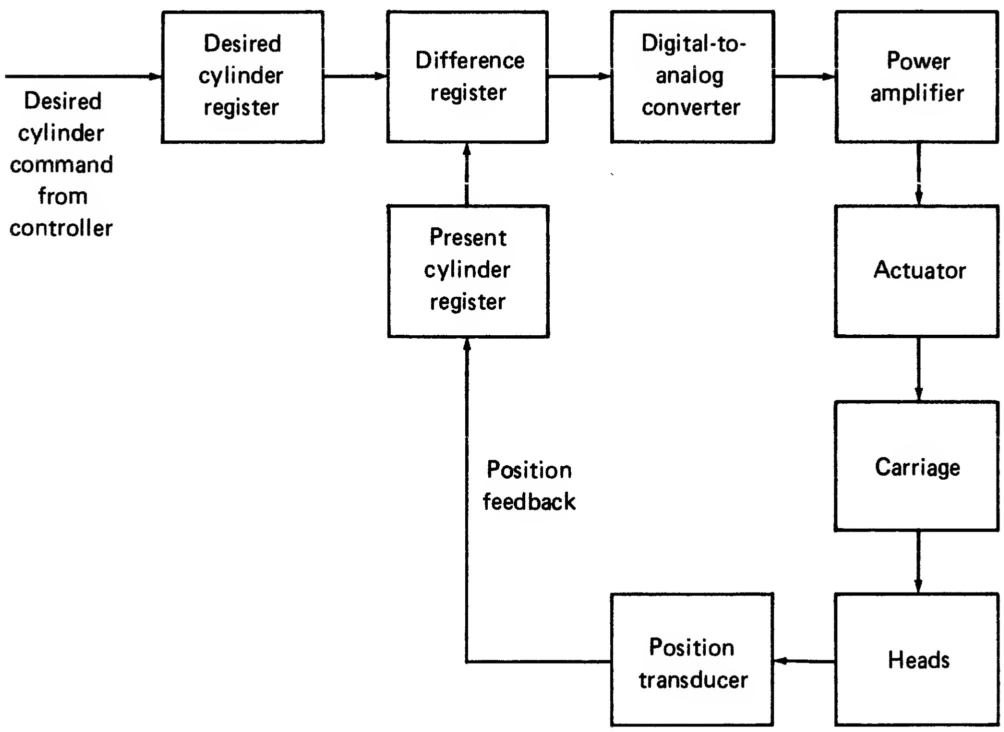

Figure 10-22. Closed loop servo block diagram (partial)

The controller instructs the disk storage device to seek a specific cylinder; the number of that cylinder is placed in a register. The desired cylinder is compared with the number of the cylinder that the heads are presently on, and the difference register output causes the digital to analog converter to generate a signal which is proportional to the distance and direction the heads must travel. This signal is amplified by the power amplifier, which applies current to the actuator; the actuator starts moving the carriage and heads toward the desired position.

As the heads move, the position transducer generates a signal indicating when the heads cross cylinders. Each time the heads cross a cylinder, the present cylinder register is updated. But as long as there is a difference between the desired cylinder and the present cylinder, the difference will continue to generate a signal which is amplified and powers the actuator. When the present cylinder is the same as the desired cylinder, the heads stop.

Unfortunately, a real closed loop servo is not quite that simple. The carriage and heads have inertia and cannot stop instantly when the heads are on the desired cylinder. You may remember from studying about velocity transducers that there is a desired velocity curve which represents the optimum velocity at all times during a seek operation. This curve allows for maximum acceleration of the carriage and heads, but decelerates them in time to stop them on the correct cylinder. Accomplishing this requires a velocity transducer (see figure 10-23).

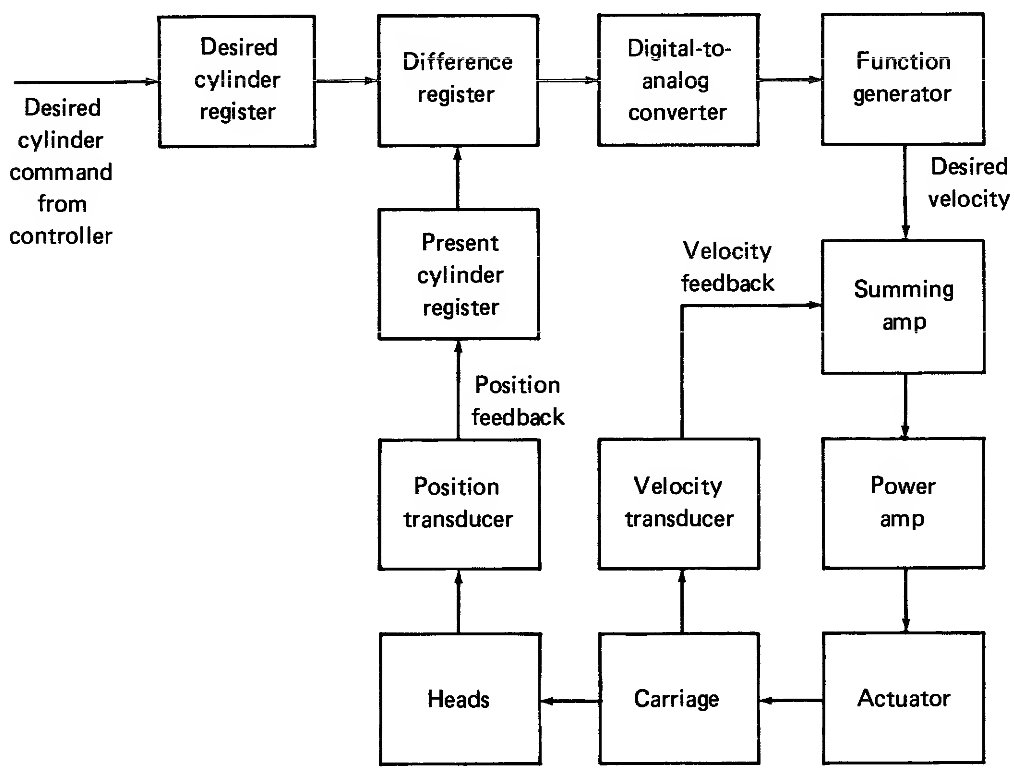

Figure 10-23. Closed loop servo with velocity transducer

In this figure 10-23, the function generator generates the desired velocity curve by modifying the linear error signal from the digital to analog converter. The summing amplifier compares the function generator signal with the actual velocity as determined by the velocity transducer. If the actual velocity is less than the desired velocity, the summing amplifier provides an accelerating signal to the power amplifier and actuator. If the actual velocity is greater than the desired velocity, the summing amplifier provides a decelerating signal to the power amplifier and actuator.

Summary

A servo is an electromechanical device that carries out instructions. There are two types of servos: open loop servos and closed loop servos. Open loop servos carry out instructions in discrete steps and do not check the results of their operations. They are simple and inexpensive, but have limited accuracy. Closed loop servos carry out instructions by constantly checking the results of their operations with the original instructions. When the results come within a specified tolerance of the instructions, the servo operation is complete. Closed loop servos are accurate and reliable, but relatively complex and expensive.

Accessing/Seeking

All disk storage devices with movable heads perform accessing, or seeking, operations. The access or seek operation (the terms are synonymous) moves the heads from one location to another on the disk pack when instructed to by the disk controller. Holding the heads on the desired cylinder is also part of the seek operation. On modern disk storage devices, most of the electronic circuitry controls seek operations.

This activity lists the phases and steps of a seek operation and describes how they are carried out by a closed loop servo.

Seek Operations

In a previous activity you learned the characteristics of open and closed loop servos. Both types of servos perform seek operations.

There are four types of seek operations. LOAD seeks occur when the device is powered up; they move the heads to the first cylinder on the pack. DIRECT seeks move the heads from cylinder to cylinder on the pack, finding addresses on instructions from the controller. Return to zero (RTZ) seeks move the heads from any point on the pack back to the first, or number zero, cylinder. RTZ seeks are usually used to clear errors in the device. Finally, UNLOAD seeks occur when the device is powered down, when power to the drive motor is interrupted, or when commanded by the control logic. UNLOAD seeks move the heads off the disk pack recording surface.

LOAD, RTZ, and UNLOAD seeks are performed in a similar manner. The heads move at a constant rate until the position feedback indicates that they are in position either at the outer edge of the disk pack or retracted. DIRECT seeks, however, are more complex.

DIRECT Seeks

DIRECT seeks are performed, generally, in two distinct phases: coarse and fine. During the coarse phase, the servo moves the heads from their present location to within one-half cylinder away from the desired cylinder. Then the fine phase takes over and moves the heads to the desired cylinder and keeps them on cylinder.

Coarse Phase

For position errors that are greater than one-half track away from the desired cylinder, the servo operates in coarse mode. In the coarse mode, the error signal is derived from

the difference register as the heads cross cylinders and the count in the difference register decreases (see figure 10-24). When the output from the digital-to-analog converter indicates that the error is greater than one-half track, the level detector switches on gate 1 and closes gate 2. When the digital-to-analog output shows that the error is less than one-half track, the level detector switches gate 1 off and turns gate 2 on.

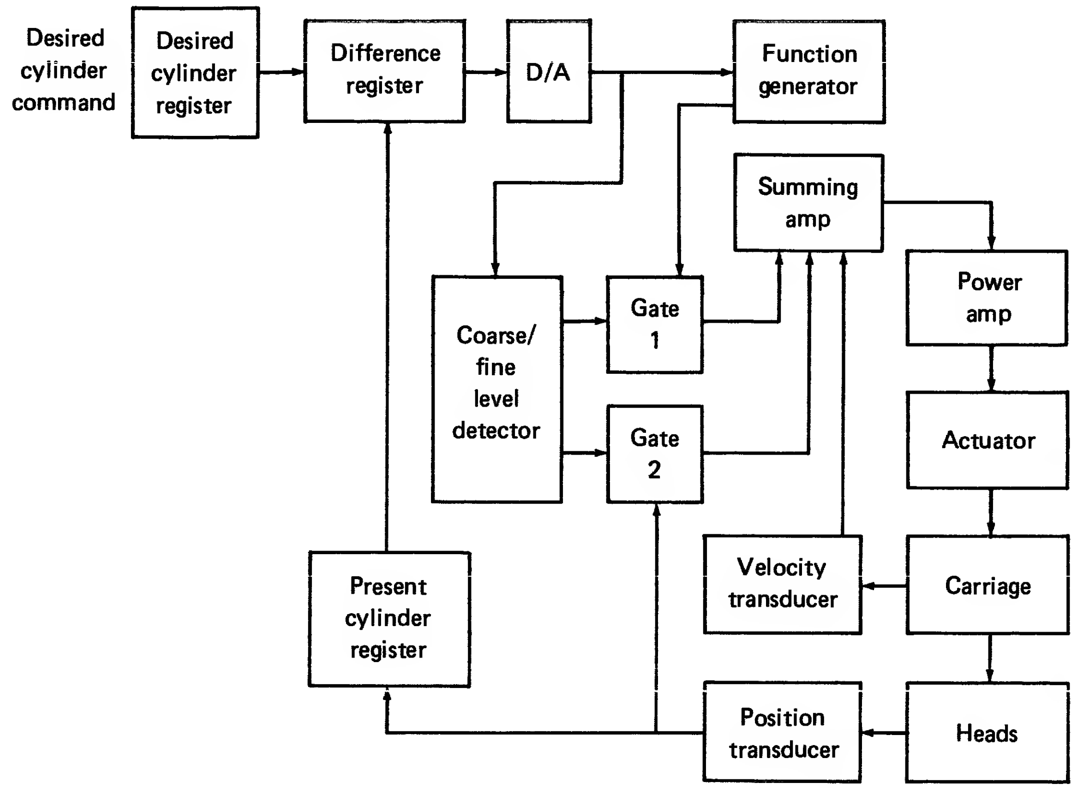

Figure 10-24. Coarse/fine servo block diagram

Depending on the length of the seek, the coarse phase generally occurs in three steps: accelerate, coast, decelerate. Figure 10-25 illustrates these three steps by showing the levels of signals generated in the servo control circuits.

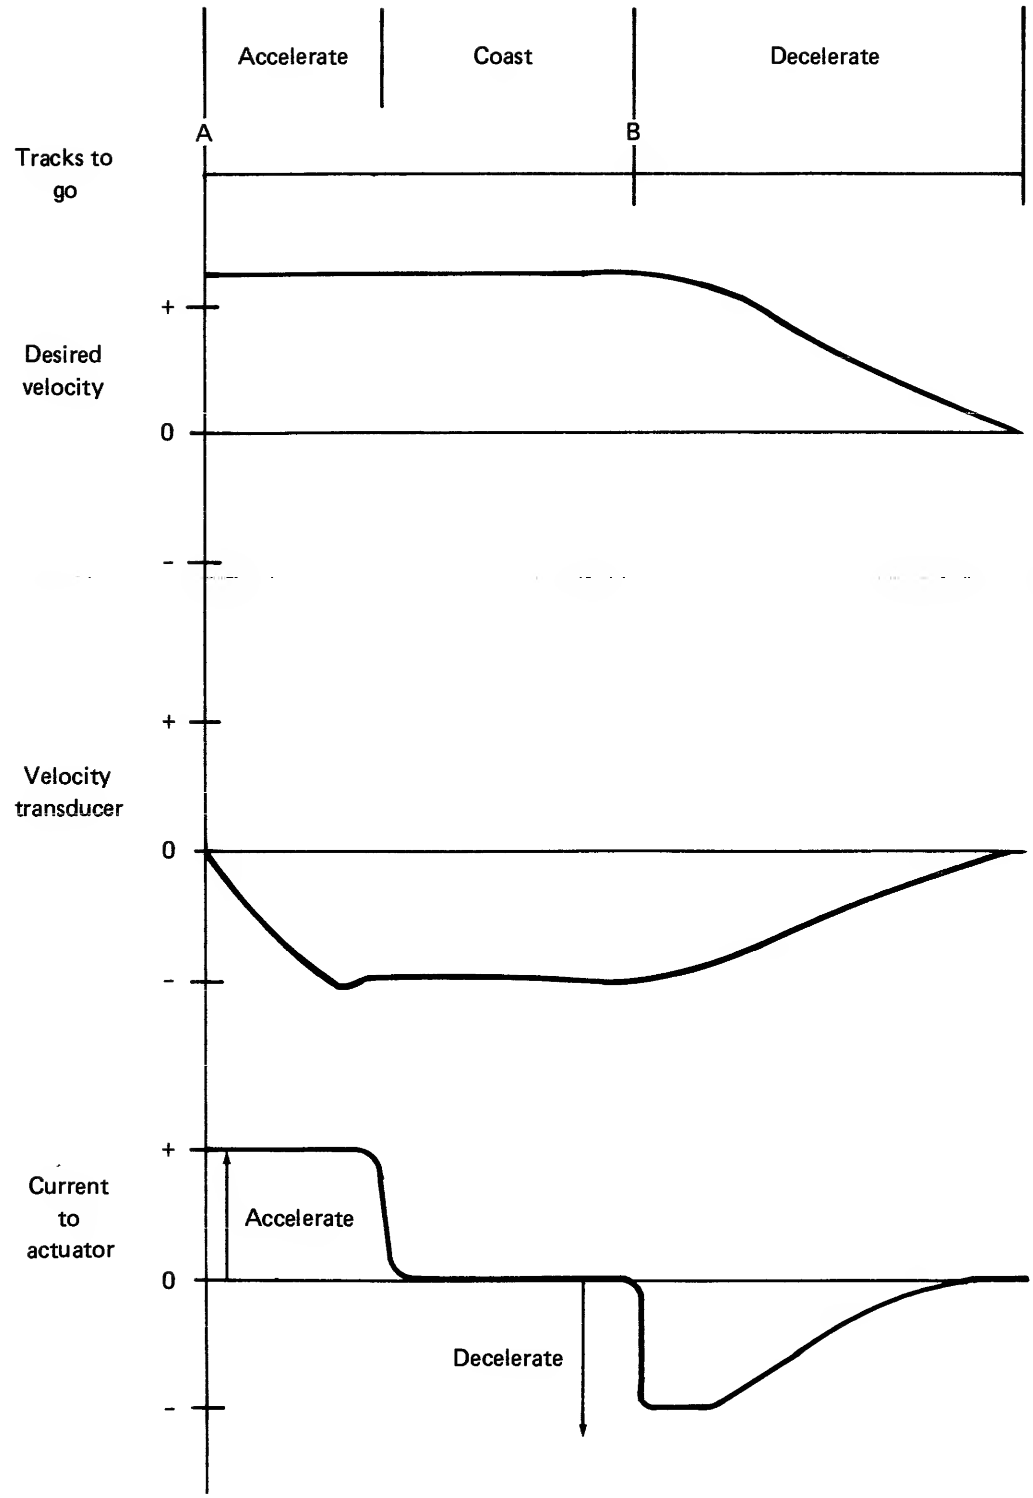

Figure 10-25. Signal levels generated in servo control circuits

"Tracks to go" is the number of cylinders remaining to be crossed to complete the seek.

"Desired velocity" is the ideal velocity curve that keeps the heads moving at optimum speed. It is generated by comparing the desired cylinder with the present cylinder.

"Velocity transducer" is the signal from the velocity transducer; it indicates the actual velocity of the heads and carriage.

"Current to actuator" is the signal applied to the power amplifier and actuator that accelerates or brakes the carriage.

Accelerate. During the accelerate step, desired veocity is clamped at a maximum level and a maximum amount of current is applied to the actuator. The carriage accelerates and when the actual velocity is equal to the desired velocity, current to the actuator is reduced to a small value which is sufficient to overcome the windage and friction.

Coast. With the current reduced, the carriage simply coasts along at a maximum speed, which is still the desired velocity. Windage and friction actually slows the carriage a little, and a small value of current applied to the actuator keeps it at a constant velocity.

Decelerate. Position feedback constantly updates tracks to go, and at point B in figure 10-25, the carriage must start to decelerate. This is similar to driving a car toward a brick wall at constant velocity; at some point you know you are going to have to start braking if you don't want to crash. The desired velocity curve represents the rate at which the heads should slow down to land right on cylinder.

When actual velocity is greater than desired velocity, a decelerating current is applied to the actuator. This reduces the actuator's velocity to the desired velocity. This deceleration continues until there is less than one-half track to go. Then the fine phase takes over.

This sequence of events represents a fairly long seek where the carriage has time to accelerate to its maximum velocity and coast. In those seeks that are less than "B long," desired velocity will start to decrease right away. The carriage will accelerate briefly and then decelerate.

Fine Phase

The fine phase of a seek operates within one-half track on either side of the desired cylinder. It uses the position feedback signal to move the heads onto the desired cylinder and keep them on cylinder during the operation. The position feedback signal around the desired cylinder is illustrated in figure 10-26.

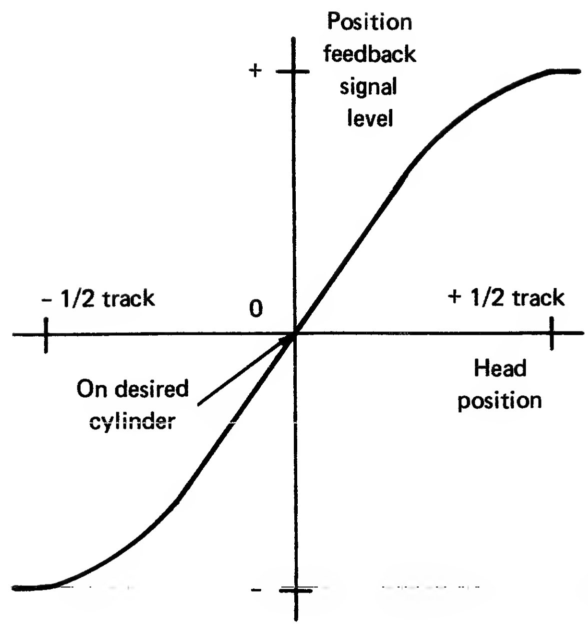

Figure 10-26. Position feedback signal

During the fine phase, the velocity feedback signal, together with the fine position signal, provides damping to the servo response. This reduces the velocity of the heads (from when the coarse phase stopped) down to zero when the heads are on cylinder.

If the heads should drift off cylinder, a position feedback signal is generated. The polarity of the position feedback signal indicates in which direction the carriage should move in order to position the heads back on cylinder. As the heads move back toward the desired cylinder, the position feedback signal decreases and the carriage decelerates to zero, putting the heads back on cylinder.

Summary

Accessing and seeking are synonymous terms for the operation that moves the carriage and heads from one location to another on the disk pack. There are four different types of seeks: LOAD, DIRECT, RTZ, and UNLOAD. Direct seeks are performed in two phases: coarse and fine. The coarse phase generally has three steps: accelerate, coast, and decelerate.

Block 11

Control Logic

Data and Control I/O—Input Signals

This learning activity presents the general concepts of data and control input/output logic. The block diagram (figure 11-1) shows the major logic areas that are common to most disk storage devices.

Data and control I/O input consists of a group of signals that are accepted by a block of receivers. The signals to the receivers are sent from the controller. From the receivers, the signals travel to the following controls and circuits: function decode, servo control, read/write control, and timing circuits.

This learning activity concentrates upon the specific input signals that are common to most disk storage devices.

Input Signals

Virtually all disk storage devices have a set of common signals. While these signals are specific in nature, the ones that will be discussed here are common to most disk storage devices. These input signals are: BUS IN, TAG LINE signals, UNIT SELECT BUS IN, UNIT SELECT TAG, WRITE DATA, and WRITE CLOCK.

BUS IN Signal

BUS IN consists of multiple lines of signals that are common to all the disks on the system that are on the same controller and are used to decode different functions. These functions include READ gate, WRITE gate, and inputs to the cylinder address register.

TAG LINE Signals

TAG LINE signals consist of the CYLINDER SELECT TAG, HEAD SELECT TAG, and CONTROL SELECT TAG.

TAG LINE signals function with the BUS IN. For example, the contents of the BUS IN signal at cylinder select time yields the information that is stored in the cylinder address register. The contents of the BUS IN signal at head select time yields the contents of the head address register. The contents of the BUS IN signal at control select time determines which of the functions should be decoded. These functions include the servo offsets, strobe offsets, READ gate, and WRITE gate.

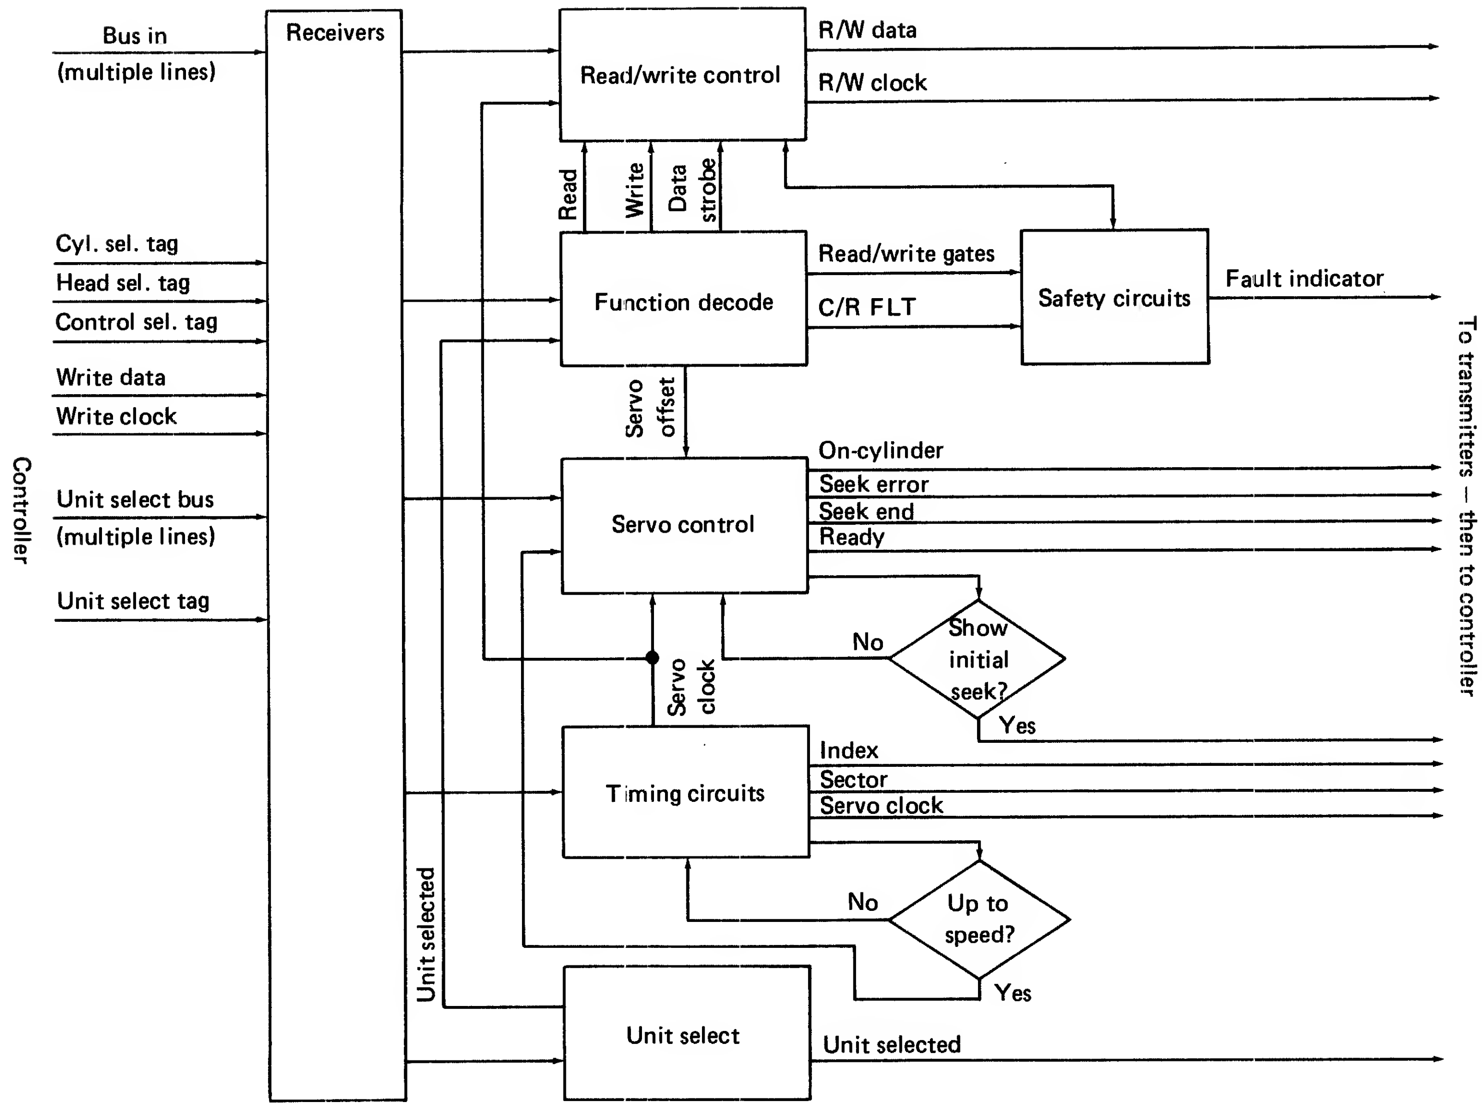

Figure 11-1. Logic areas common to disk storage devices

UNIT SELECT BUS IN and UNIT SELECT TAG Signals

The UNIT SELECT BUS IN signal also consists of multiple lines. The contents of the UNIT SELECT BUS IN at UNIT SELECT TAG time determines the unit number that will be selected. If that signal compares favorably with the unit number which is permanently assigned to that particular machine, it generates a UNIT SELECT signal. This signal essentially enables the entire I/O device, both receivers and transmitters. If you do not get a unit select compare, the function decode is inhibited from interpreting any command functions.

WRITE DATA Signal

This signal carries NRZ data that is going to be recorded on the disk pack.

WRITE CLOCK Signal

This signal is synchronized to the NRZ write data, and it is a return of the servo clock. This signal is transmitted continuously from the controller.

Data and Control I/O—Output Signals

This learning activity presents the data control output signals that are common to most disk storage systems.

Data control output signals are sent to transmitters and, in turn, to the controller. These signals are sent from the function decode, servo control, read/write control, timing circuits, safety circuits and unit select. The block diagram (figure 11-2) shows the major logic areas that are common to most disk storage devices.

Output Signals

Virtually all disk storage devices have a set of common signals. While these signals are specific in nature, the ones that will be discussed here are common to most disk storage devices. These output signals are: ON-CYLINDER, SEEK ERROR, SEEK END, READY, FAULT, INDEX, SECTOR, SERVO CLOCK, READ DATA and READ CLOCK.

ON-CYLINDER Signal

The ON-CYLINDER signal means that the positioner is on track and that the unit can accept commands for reading or writing. This signal comes from the area of servo control.

SEEK ERROR Signal

The SEEK ERROR signal means that the disk drive has not successfully completed a seek operation, and that it has exceeded the legal seek boundaries in some manner. This signal comes from the area of servo control.

SEEK END Signal

The SEEK END signal is usually the same as ON-CYLINDER and differs only when you get a SEEK ERROR signal. When you receive a SEEK END signal, you should sample ON-CYLINDER and SEEK ERROR to judge what condition exists. This signal comes from the area of servo control.

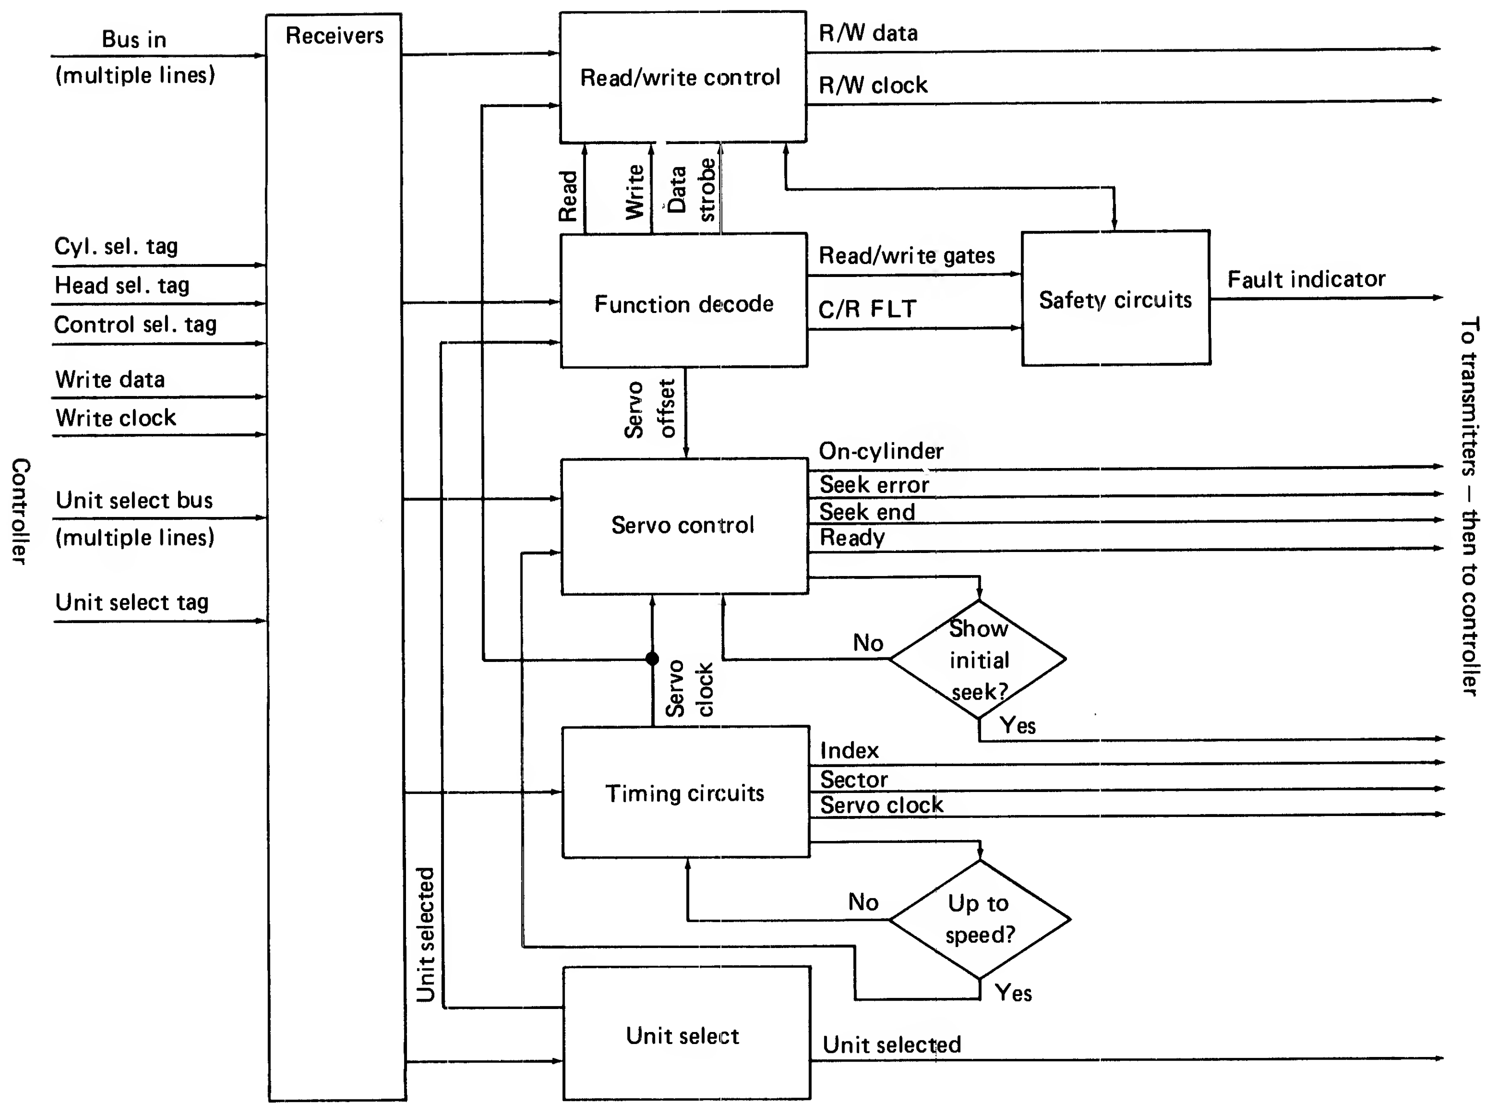

Figure 11-2. Logic areas common to disk storage devices

READY Signal

The READY signal means that the spindle is up to speed and that an initial seek has been completed. It also means that the unit is positioned at a cylinder, which is usually cylinder 0, and that it can accept a command for reading or writing. However, the unit must be up to speed and must have completed an initial seek before the READY signal will activate. This signal comes from the area of servo control.

FAULT Signal

The FAULT signal comes from the safety circuits and means that an error condition exists in the machine.

INDEX Signal

The INDEX signal is a reference signal that indicates the start of a track. One INDEX pulse is generated for each revolution of the disk. This signal originates from the timing circuits.

SECTOR Signal

The SECTOR signal is divided into multiples because there is usually more than one sector per track. A sector is a subdivision of the tracks, so there may be 16, 32, 64, 128 (etc.) numbers of sectors between index pulses. This signal originates from the timing circuits.

SERVO CLOCK Signal

This signal is the clock that runs at the frequency at which you are attempting to read or write. In most devices, this signal is sent to the controller and the controller sends it back as a WRITE CLOCK signal. In the controller, the SERVO CLOCK signal is designed to be in sync with the WRITE DATA, so that when it gets to the drive it has the proper polarity and will clock correctly. This signal originates from the timing circuits.

READ DATA, READ CLOCK Signals

READ DATA, and READ CLOCK are signals that come from read/write control. These signals should be self-explanatory.

Powering-Up Sequence

This learning activity discusses the sequence for powering on a storage module drive. It also covers cabling, the sequence for powering off, and emergency retract.

Power-Up Sequence

A storage module drive can be powered up from two modes: local and remote. Local mode refers to a condition in which the drive is off-line. Remote mode exists when the storage module drive's power-on and power-off sequence is under the control of the controller, or of another storage module drive. This learning activity will omit local mode power sequencing and refer only to the power sequencing that is initiated by the controller. (See figure 11-3.)

The block diagram in figure 11-3 refers to the power-on sequence for one storage module drive in the remote mode. Although some areas of this block diagram are self-explanatory, they will all be briefly explained.

A successful power-up sequence is based on two conditions. The first condition is that the ON-OFF switch is in the on position. The second condition is that the disk storage device that is being powered up receives a PICK AND HOLD signal from the controller or from another drive. The PICK AND HOLD signal, which is sent by the controller, can also be referred to as the POWER SEQUENCE signal and it provides the power-up and power-down sequence for the drives. If the ON-OFF switch is in the on position, then the start relay will energize when power becomes available. The power sequence will continue only if the PICK AND HOLD signal is received from the controller. This sets the remote start flip-flop (FF). A SEQUENCE HOLD signal from the controller holds the remote start FF in the set state, and this maintains the start drive motor signal.

The drive power-on sequence starts the drive motor and places the head in the home position. This sequence is initiated by closing the ON-OFF switch at the rear of a drive unit. This energizes the sequence relay.

When the motor triac is initially picked, a fixed delay starts which allows the spindle to get up to speed. If the spindle fails to come up to speed within that allotted time, the motor pick drops, the brake is applied, and the drive motor shuts down. Once the spindle speed reaches 75 percent of its peak speed, the logic to load the head is enabled. After a fixed delay, the heads are loaded to track 0. Once the spindle is up to speed, the speed relay energizes. The power supply connects the power amplifiers to the voice coil on the actuator and connects the emergency retract capacitor. This prepares the voice coil to respond to commands from the servo logic and charges the emergency retract capacitor so that it is ready for the emergency retract condition. In addition, a PICK signal is sent out to the next drive or to the controller.

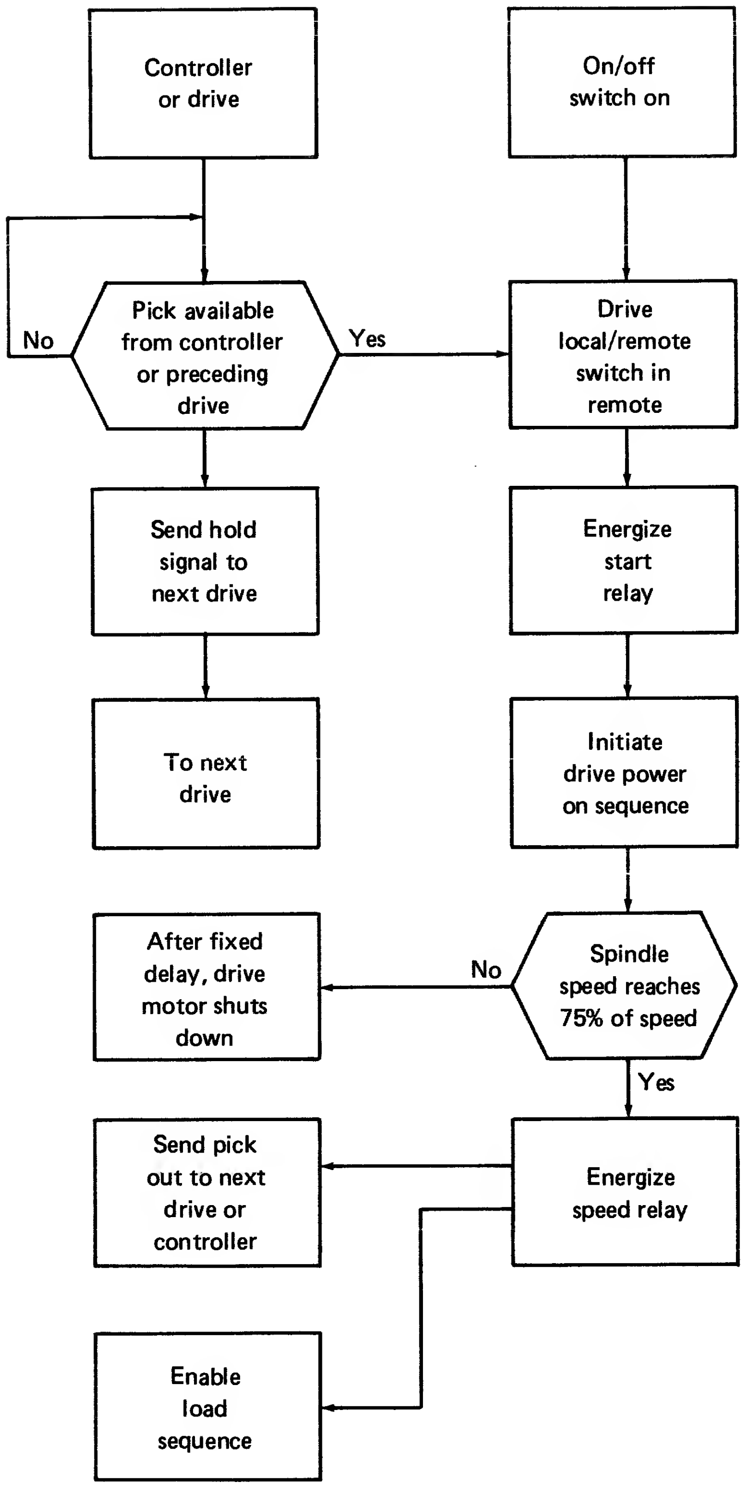

Figure 11-3. Power sequencing

Cabling

Most disk drives have what is called a two-cable interface (see figure 11-4). This refers to the two major types of cables, a unique cable, and a daisy chain cable. A unique cable is a direct link between the controller and the individual drives in the system. The daisy chain cable is a link that runs from the individual units to the controller via other storage module drives in the system. The unique cables run directly from the controller to the individual disk storage devices and have no other units between themselves and the controller. The path along a daisy chain cable goes from the controller to drive unit one, then to drive unit two from drive unit one, and so forth.

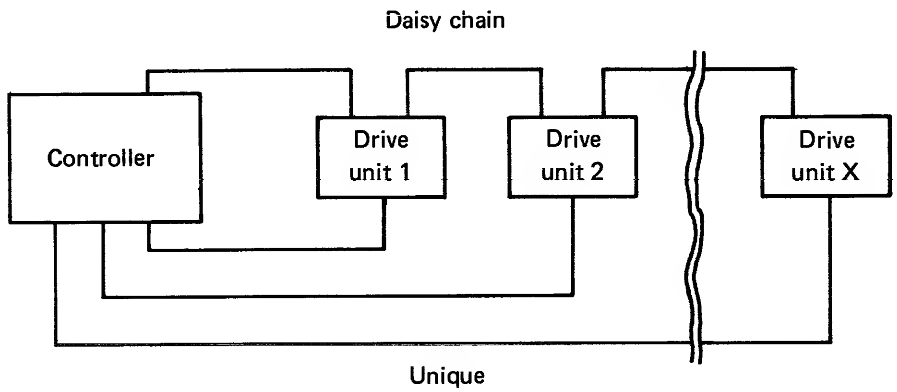

Figure 11-4. Cabling configurations

The unique cable generally carries the READ/WRITE signals. A daisy chain cable usually carries all of the TAG LINE signals and BUS IN lines. In the output area, the unique cable carries the READ DATA, READ CLOCK, SERVO CLOCK, and UNIT SELECTED. The others are located in the daisy chain cable.

Some storage module drives only have one type of cable. However, the vast majority of storage module drives have a two-cable configuration.

Power-Off Sequence

The power-off sequence unloads the heads and stops the drive motor. The power-off sequence begins when the ON-OFF switch is opened. This activates the retract circuit and retracts the heads.

Opening the ON-OFF switch removes power from the brake coil and enables the emergency retract circuit. This circuit retracts the heads and holds the actuator in the retract position. If the speed of the spindle drops below 75 percent of its normal speed, the retract circuit activates and the drive motor brake is applied.

Emergency Retract

The emergency retract function retracts the head from the data area if the disk speed is reduced to a dangerously low level. If the speed falls below 75 percent of its normal speed, and the heads are not retracted, it could result in a head crash and/or in the loss of recorded data.

Emergency retract will usually activate under the following conditions:

- Loss of AC power. If AC power is lost, then DC power is also lost and an emergency retract occurs.
- Loss of speed. If the spindle motor drops below 75 percent of its normal speed, a speed detection circuit causes an emergency retract to occur.
- Loss of sufficient voltage. This causes a braking action on the motor and an emergency retract occurs.
- Drive motor thermal overload. If a drive motor overheats, a circuit-breaker opens within the motor. This cuts off power to the motor, which slows down; the loss of speed causes an emergency retract to occur.

Function Decoder

This learning activity presents the function decoder. It explains how the function decoder changes signals and what those signals become.

The function decoder consists of AND gates. The function decoder contains the contents of the bus at control select time. It determines the functions that are going to be performed. It does so by placing an AND on each of the bits on the bus with the control select tag.

A disk storage device usually utilizes a 10-bit bus. At control select tag time, different signals are decoded. These signals are used to initiate various operations that will be performed by the disk storage device (see figure 11-5). The contents of these bus bit lines are as follows:

- Bit 0 yields a WRITE gate, which enables the write drivers.
- Bit 1 yields a READ gate, which enables the digital read data lines.
- Bit 2 yields a servo offset positive, which offsets the actuator from the nominal on the cylinder positioned toward the spindle.
- Bit 3 yields a servo offset negative, which offsets the actuator from the nominal on the cylinder positioned away from the spindle.
- Bit 4 yields a controller fault clear, which is a 100 ns (minimum) pulse that is sent to the disk storage device. This pulse will clear the fault latch if a fault condition no longer exists.
- Bit 5 yields an address mark enabled, and when this is combined with a WRITE gate, the address mark is written. When it is combined with a READ gate, an address mark search is initiated.
- Bit 6 yields an RTZ seek, which is a pulse that is sent to the drive which causes the actuator to seek to track 0. It also clears the head address register and clears a seek error latch.
- Bit 7 yields a data strobe early, which enables the phase lock oscillator (PLO) data separator to strobe the data earlier than optimum time.
- Bit 8 yields a data strobe late, which enables the PLO data separator to strobe the data later than optimum time.
- Bit 9 yields a release, which is only applicable to dual channel units. This line resets the reserve and disable flip-flop for both channels, and releases both the reserve and priority select commands.

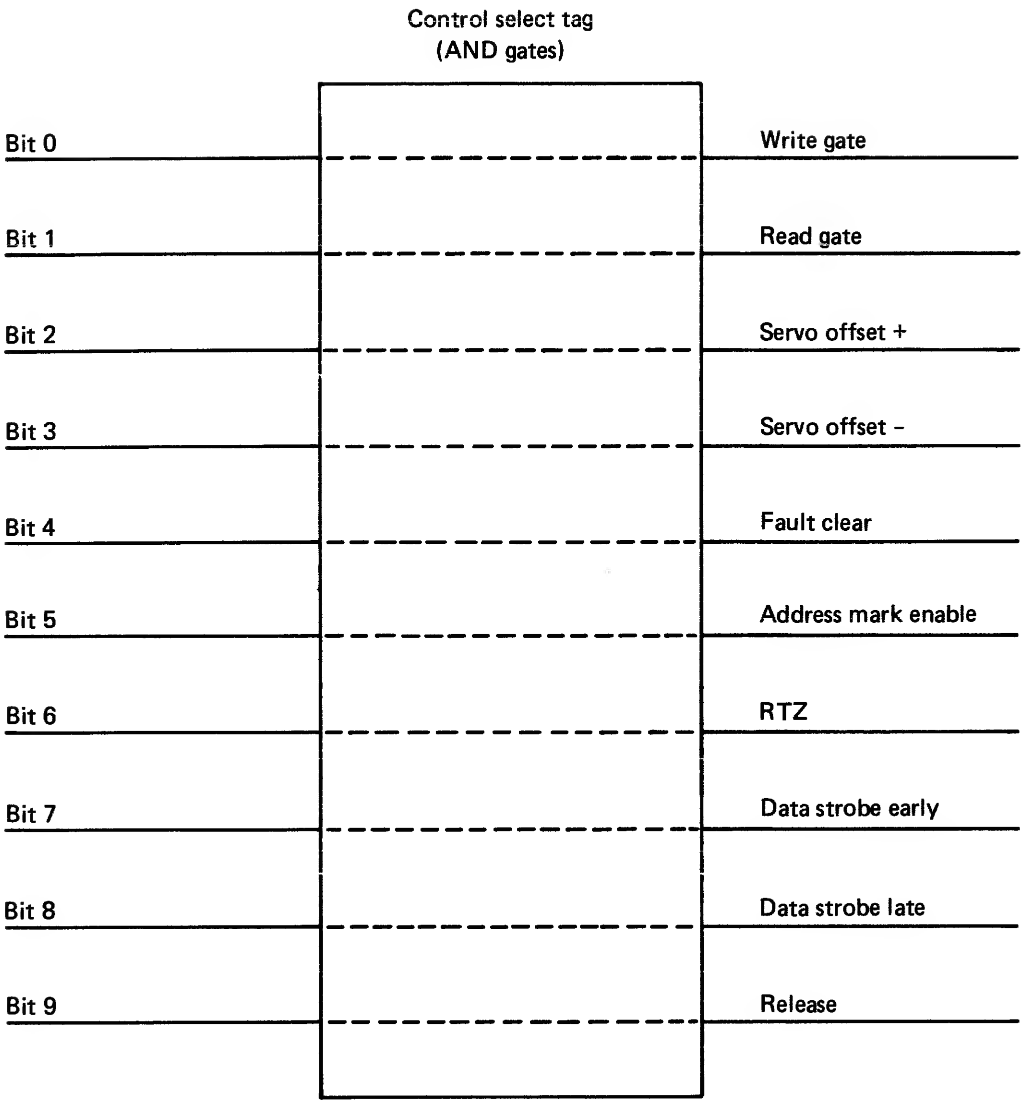

Figure 11-5. Function decoder

Servo Control

This learning activity discusses, in some detail, the area of logic known as servo control. The major areas that will be covered include an overview of servo control logic, the digital-to-analog converter, the summing amplifier, modes and seeks, the output of the servo head, the microprocessor servo, and the track servo signal. See figure 11-6 for a diagram of servo control.

Servo Control Logic

BUS IN and CYLINDER SELECT signals are sent by the receivers to the servo control area. In servo control, the contents of the bus at cylinder select tag time are loaded into the cylinder address register, (CAR). The contents of this register are compared with the contents of another register, called the present address register.

The present address register contains the binary coded present address. When a new desired address is loaded into the CAR, the two addresses are compared. This comparison takes place in the subtractor. The difference is then loaded into the difference counter. This difference is a measure of the seek length that will be attempted. The difference counter is strobed in by cylinder select, and a delay is built in from the time it strobes into the present address. The delay is necessary in order for the subtractor to work. In most cases, the cylinder address register is strobed on the leading edge of the pulse, and the difference counter is strobed on the trailing edge of the pulse. These pulses are usually one microsecond long.

The difference counter is then decoded—there are many types of signals that may be used for this. In figure 11-6, the equation T = 0 is drawn. This is a common equation that is used in such places as on-cylinder detection.

The most important function of the difference counter is that it provides an input to the digital-to-analog converter. This converts the difference information from a digital signal to an analog voltage. This voltage is a measure of how far you have to go when you start a seek.

Digital-to-Analog Converter

The signal that comes out of the digital-to-analog converter is an analog voltage that is proportional to the magnitude of the binary count in the difference counter. As cylinders are crossed, the cylinder count goes down. Initially, the difference counter has a high count, and, as the seek comes closer to its final destination, it keeps crossing cylinder pulses and the binary count goes down. Therefore, the analog voltage that

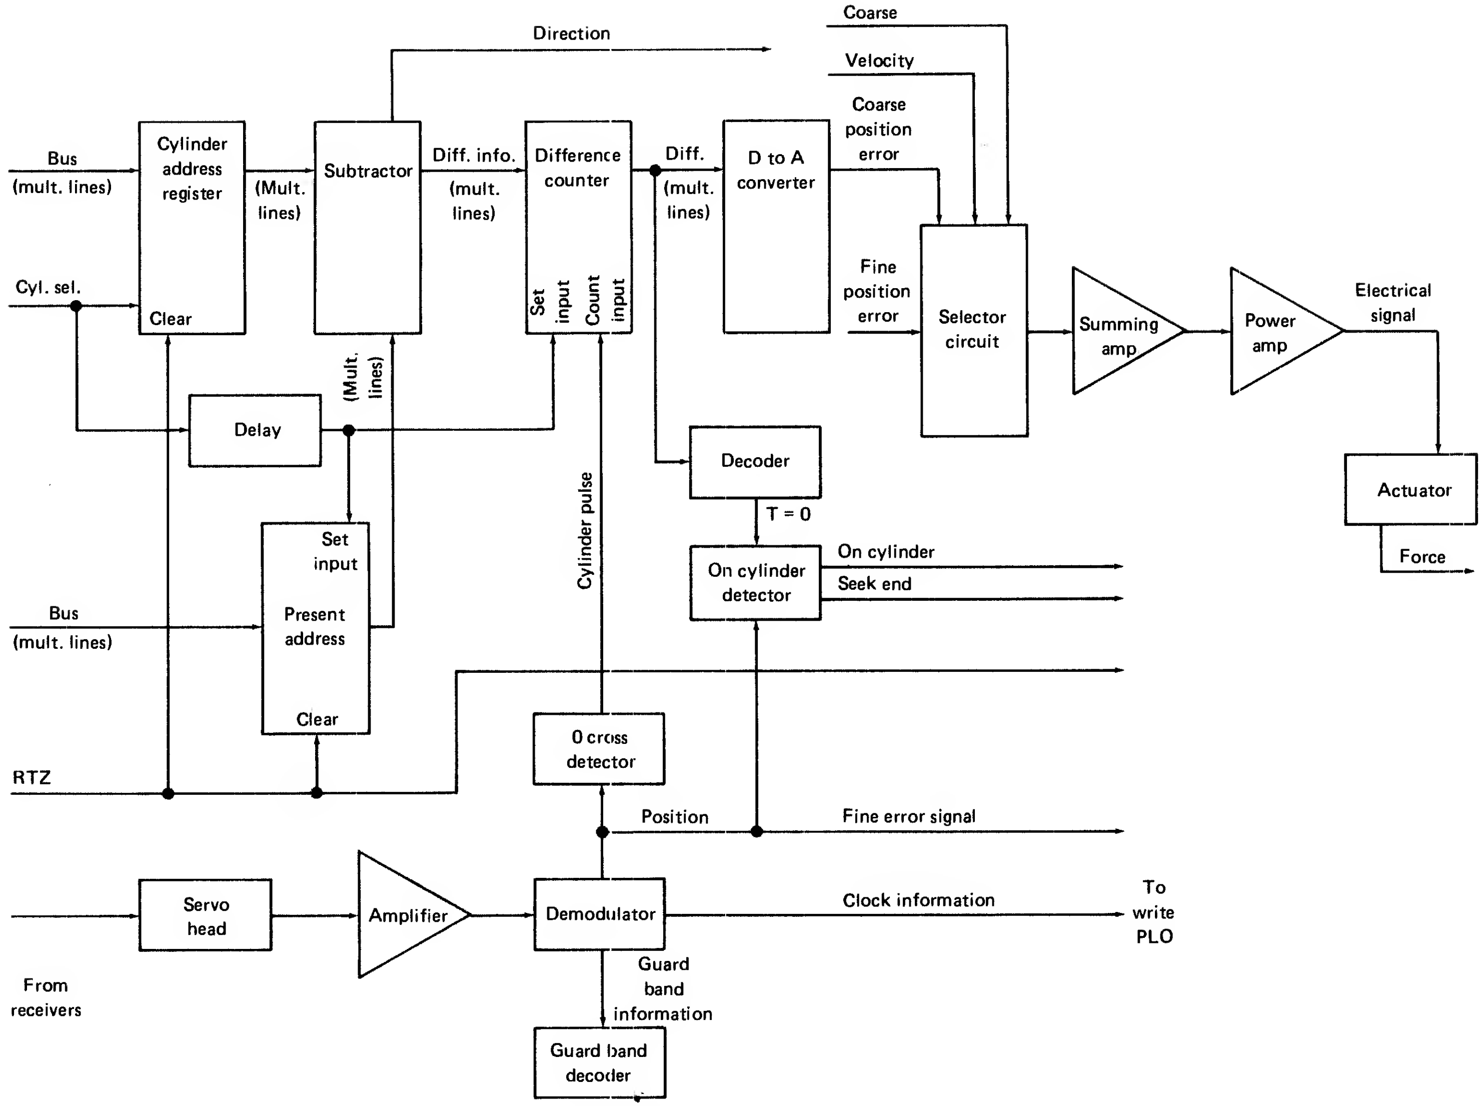

Figure 11-6. Servo control

comes out of the D-to-A converter also goes down. This means that the error signal is continuously getting smaller. It keeps getting smaller until it is on-track and it goes to zero.

A digital-to-analog converter counts down one for each time a track is crossed. This is why the signal appears like a stairway (see figure 11-7).

The smoothed curve in figure 11-7 refers to the desired velocity or coarse position. The sum of velocity is always distance. Therefore, if you add the velocity between track pulses, you receive a voltage that is proportional to distance. By summing, or inverting those two factors, a desired velocity or COARSE POSITION ERROR is received. This is the error signal that the servo is attempting to reduce.

Summing Amplifier

~~There are many things that occur within a summing amplifier. What is most important~~ is that the coarse position error and the velocity are usually summed at that amplifier. At that point, the total summed at the amplifier goes to another amplifier which is called a power amplifier. Then an electric signal is sent to the actuator; then, force, or a mechanical motion, comes out of the actuator.

A selector is located between the summing amplifier and the signal that comes from the digital-to-analog converter. As previously stated, the signal that comes from the digital-to-analog converter is called the COARSE POSITION ERROR. However, both FINE POSITION ERROR and COARSE POSITION ERROR signals are sent to the selector circuit, depending upon the mode that the machine is operating in.

Mode and Seek Conditions

If the machine is operating in a coarse mode, it enables a COARSE POSITION ERROR signal into the summing amplifier. If the machine is operating in a fine mode, it sends a FINE POSITION ERROR signal into the summing amplifier. The machine must always be operating in either one of these modes, unless it is in an initial seek condition, or a return to zero seek condition.

These two types of seeks will not be discussed in detail in this learning activity. This is because these seeks are almost never the same on any two machines. Therefore, they cannot be discussed in a general or universal manner. However, one thing that can be generalized is that a return to zero, or load command, defeats some of the normal seeking circuitry and accomplishes a positioning to cylinder Ø through separate circuits.

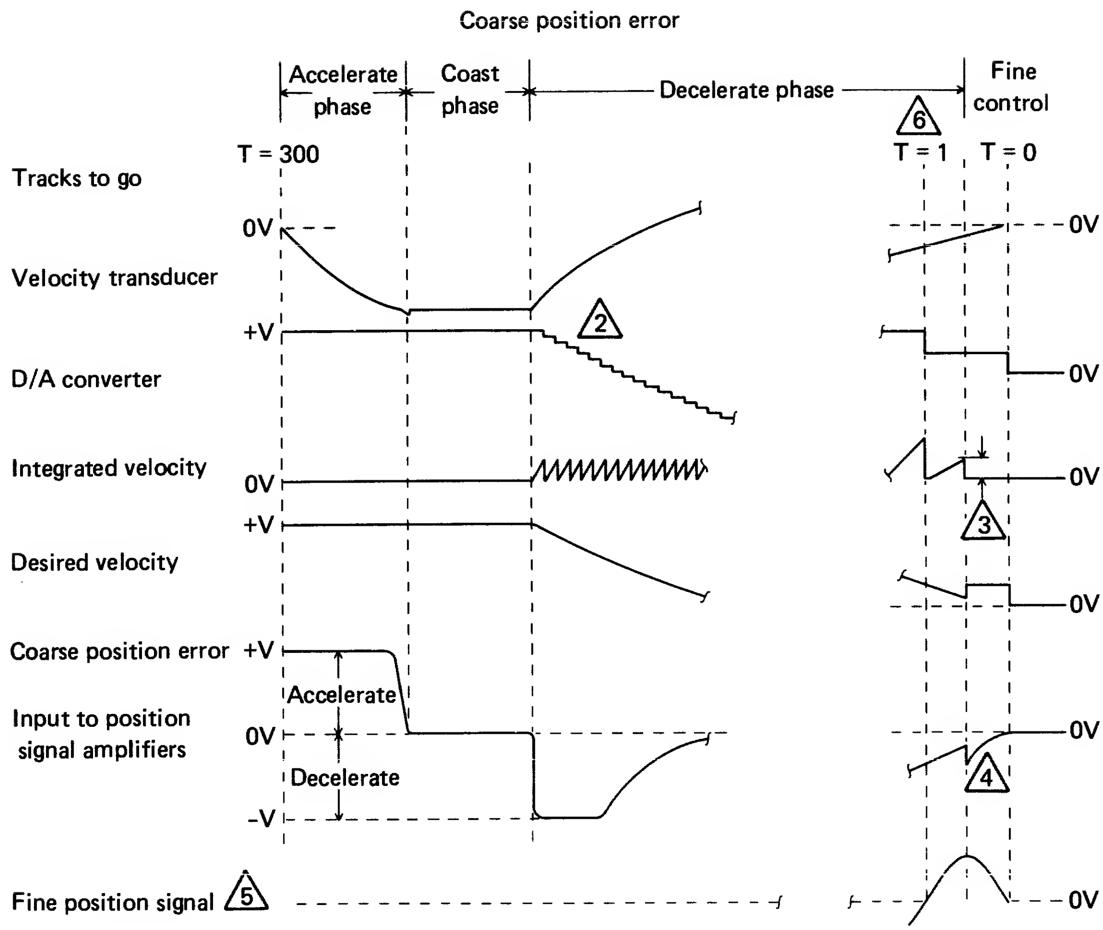

NOTES:

1. Signals shown apply to FWD seek about 300 cyl in length. All polarities except D/A converter are opposite for rev seeks. Timing and amplitude are not to scale.

2. Output decreases with each cylinder pulse.

3. Servo system switches to fine control when integrated velocity exceeds 0.9V.

4. Gain change caused by switch from coarse to fine control. Signal for last half track is due to file position input and is shown here for reference only.

5. From fine position control circuits and is shown for reference only.

6. Scale expanded for clarity beyond T = 1.

Figure 11-7. Direct seek coarse position control signals

Servo Head

The output of the servo head is amplified and sent into a demodulator. The demodulator obtains the position information by decoding the servo head information.

Guard band information also comes out of the demodulator. This is encoded in areas that are outside the legal track boundaries. At all times, the information is on one side of the boundaries. Most machines utilize guard band information to inform you what band it is in. This information is also used to enable the machine to get back to zero. The majority of disk storage devices initiate an automatic return to zero whenever they get outside the legal track boundaries. In other words, they function similarly to a safety circuit.

Microprocessor Servo

The microprocessor servo is a relatively new feature. What the microprocessor servo does is replace the area that is outlined on figure 11-8 with the microprocessor and its associated hardware. All of this logic, including the D-to-A converter, is replaced by the microprocessor. Microprocessor servos will obtain a larger and larger share of the disk storage device market.

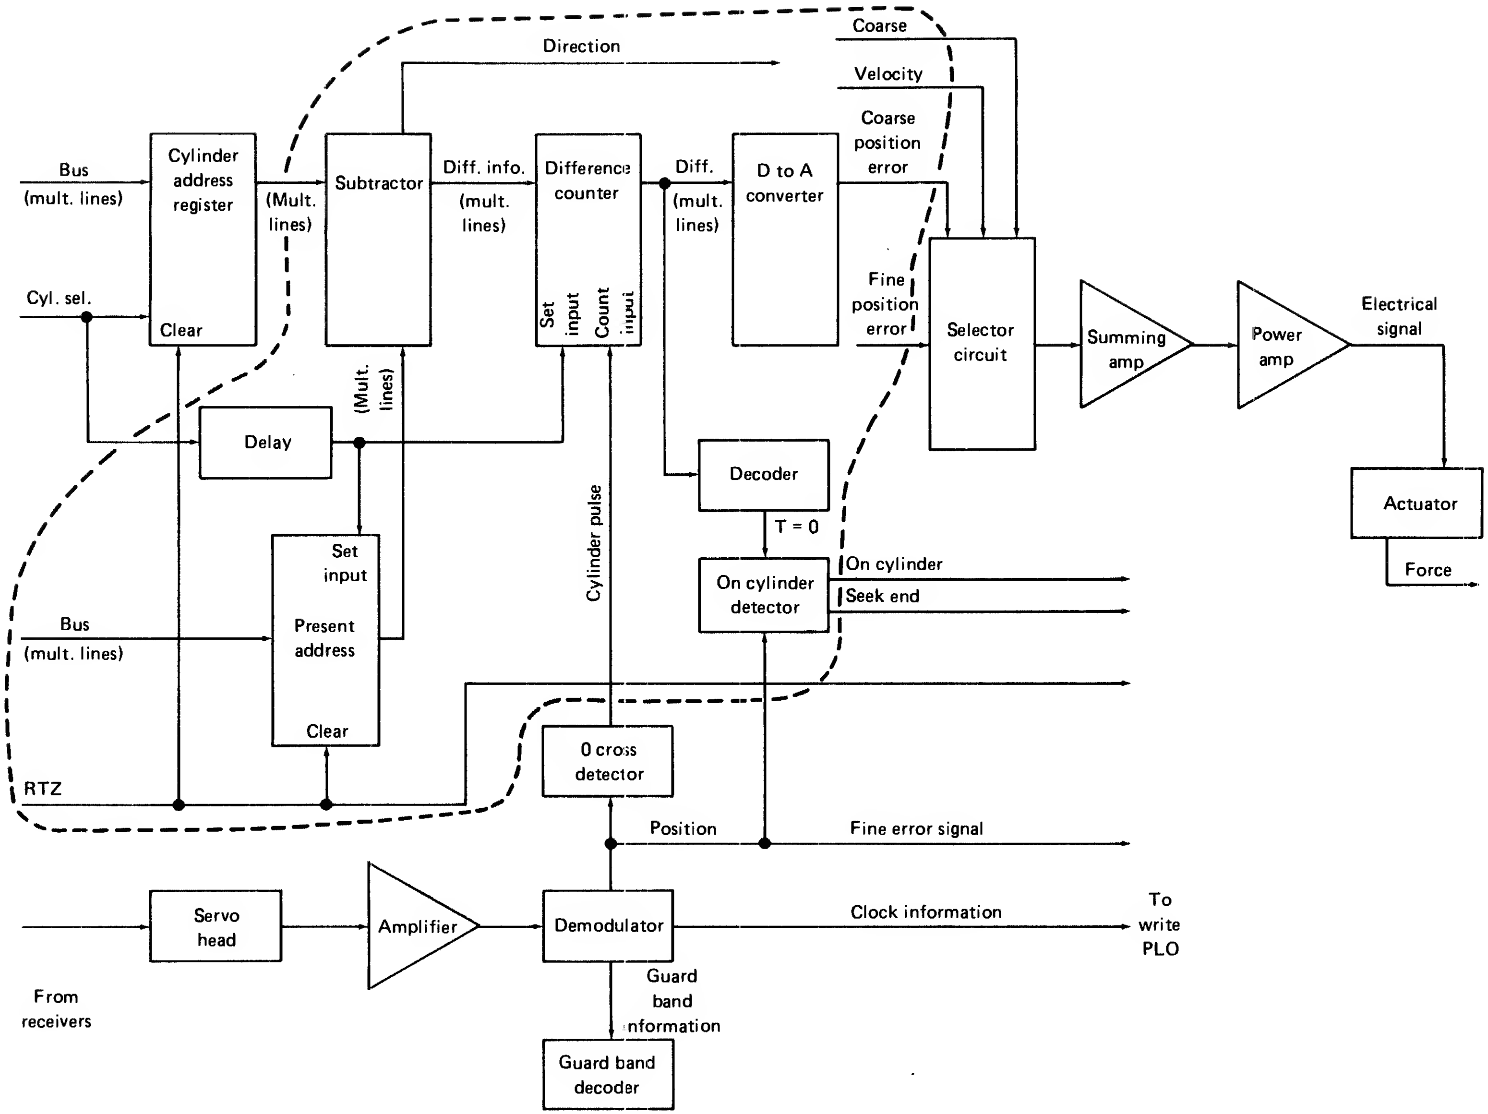

Figure 11-8. Area replaced by microprocessor servo control

Read/Write Logic

This learning activity discusses, in some detail, the logic for read/write operations. Read and write operations and a brief discussion on peak shift are provided.

Read Operations

See figure 11-9 for a flowchart of read/write logic.

For read operations, the desired head number goes from the receiver into a head register. The head register is a series of flip-flops and the input is loaded into the register at head select tag time. When these signals come out of the register, they go into the decoder. In some machines, the signals go directly from the head register to the heads. However, it is more common for the signals to go through a decoder and a translator, before they travel to the heads.

There are usually multiple arms on the head register and on the decoder. These arms then travel to a level shifter of some type which is called the translator or invertor. From the translator, they travel directly to the heads.

When the signal goes from the translator to the heads, the line generally goes to the center tap of the heads, sets it at a predetermined voltage, and selects the correct head. Generally, there is a diode matrix located in the heads.

The input that comes from the heads goes directly to the amplifier. From the amplifier it goes to a filter and then to an analog-to-digital converter. The analog-to-digital converter changes the analog signal to a digital voltage. From this converter, the signals move into a data latch and then to a read phase lock oscillator (PLO) and data separator. NRZ read data and read clock signals are sent from the read PLO and data separator. These signals go to the transmitters, which ultimately go to the controller.

The function decode sends a READ gate signal to the read PLO. There is also a flow from the write driver into the heads. When the machine is reading, it shuts off the write driver and flows into the amplifiers.

Data latch is a circuit that is designed to prevent errors that result from over-resolved signals. It also helps in recovering data which may be damaged due to a media defect.

The analog-to-digital converter, and the read circuit function, changes the analog MFM data into digital MFM data. This data is sent to the read PLO and data separator.

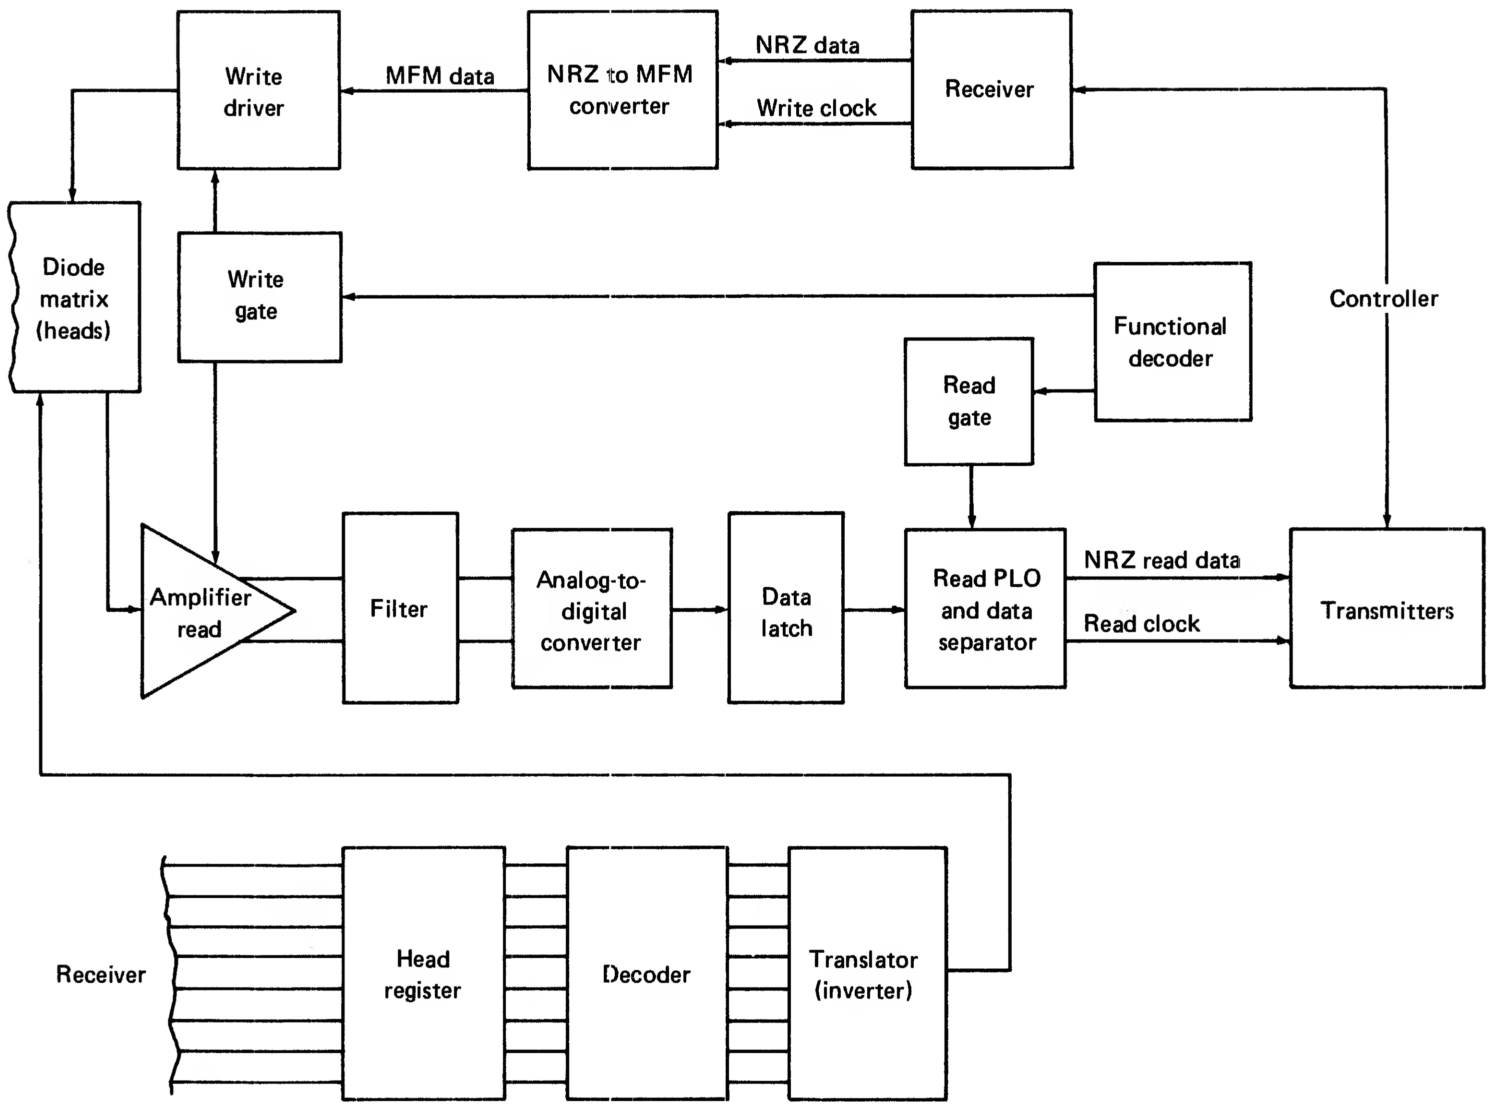

Figure 11-9. Read/write logic

The read PLO takes the peak shift out of the data. In addition, the oscillator in the PLO has a high amount of inertia. When data pulses enter the read phase lock oscillator from the data latch, that oscillator adjusts its frequency to coincide with the data pulses. This puts the timing in sync with the revolution time of the spindle. If the spindle slows down, the pulses that come into the read PLO will slow down and the oscillator in the PLO will also slow down. Systems that have a read PLO can withstand a wider range of spindle speeds.

The read PLO and data separator have two basic functions. First, they convert the MFM data from the analog-to-digital converter into NRZ data. Second, they generate a READ CLOCK signal, which is locked into the frequency of the read data. Both the NRZ data and the READ CLOCK signal are transmitted to the controller via transmitters.

The data separator circuit determines if the data pulses represent a 1 or a 0, and then converts this data to NRZ. The data separator circuit also generates a feedback clock pulse to the comparator and the READ CLOCK signal. The feedback clock pulse is eventually sent to the controller.

Write Operations

For write operations, input is received from the controller and travels to the receiver. The receiver sends NRZ data, and a WRITE CLOCK signal, into the NRZ to MFM converter. This converter receives input in the form of NRZ data and converts this to MFM data. After conversion, the data is sent to the write drive and on to the heads.

For both reading and writing operations, the function decode sends a signal to a gate. For write operations, the function decode sends a signal to the WRITE gate to enable the write driver. In addition, it also sends a signal to the read amplifier and disables it when the machine is writing. When the machine is reading, signals that turn on the amplifier and disable the write driver are sent. In order to fully understand the functions of the NRZ to MFM converter, it is necessary to understand MFM data and peak shift.

Assuming that you have some knowledge of MFM data, you should know that there are some patterns that need to be shifted from nominal in order to improve the margin when you read the signal back. Those patterns are called write compensation timing (see figure 11-10). In this illustration, the 100 pattern is a pulse that gets shifted late, 011 also gets shifted late, 001 gets shifted early, and 10 gets shifted early. The other patterns that are shown in this illustration are not effective.

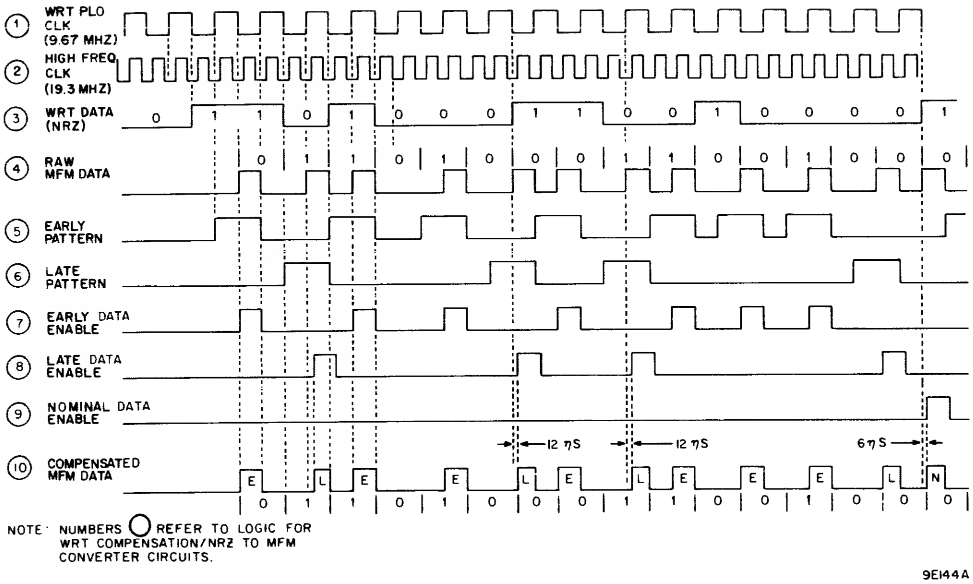

Figure 11-10. Write compensation timing

An important event will be taking place here. This event is called peak shift. (See figure 11-11.)

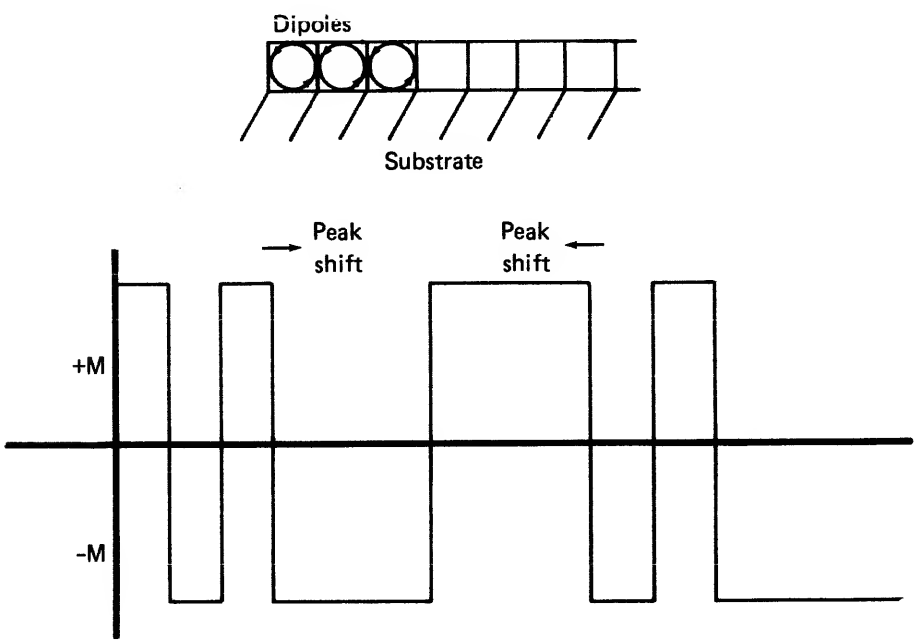

Figure 11-11. Peak shift

When the machine is writing data on the disk, there are going to be areas of high-density pulses and low-density pulses. When a high-frequency pulse is written next to a low-frequency pulse, the pulse crowding tends to shift the high-density transition toward the low-density transition. Transition is where the magnetization of the dipoles changes from a positive to a negative state. This is what is called peak shift.

To correct this condition, the machine attempts to compensate for it during the write operation. In other words, it attempts to shift it backwards and spread it in a wider than normal distance. This is called write precompensation. Then when the data peaks shift and become narrower, the result should be that the data peaks will return toward nominal, or normal. In most cases, write precompensation is not 100 percent effective; however, it does help a great deal in correcting for peak shift.

In writing operations, the write driver takes the data and provides the current source to the heads. In essence, the write driver toggles the head with the correct amount of current. The most important thing in the write driver is that it has the correct current magnitude. If the current is too high, the head will saturate and it will not produce a high-quality signal on the disk. If the current magnitude is too low, then the disk will not be saturated and the old data will not be erased.

Fault Indications

This learning activity lists the most common conditions that are interpreted as errors by the disk storage device. All of these conditions either light an indicator on the device, or send a signal to the controller indicating that an error condition exists. These errors are divided into two categories: 1) the errors that are indicated by a fault latch, and 2) the errors that are not indicated by a fault latch.

Errors Indicated by a Fault Latch

The errors that are indicated by a fault latch are:

- General
- Write fault
- Head select fault
- Read and write fault
- (Read or write) off-cylinder fault
- Voltage fault

General

General errors describe certain errors that set the drive's fault latches when error conditions are recognized. Setting the fault latch accomplishes four things:

- It enables the fault line to the controller.
- It lights the fault indicator on the drive's control panel.
- It clears the drive unit READY signal.
- It inhibits the drive's write and load circuitry.

All of these events prevent any further drive operations from being performed until the error is corrected and the fault latch is cleared.

After the error is corrected and an error condition no longer exists, the fault latch can be cleared by any of the following operations:

- Pressing the fault switch on the operator's panel
- Obtaining the controller FAULT-CLEAR signal from the controller
- Pressing the maintenance FAULT-CLEAR switch on the fault card
- Powering down the unit

Whenever an error occurs that sets the fault latch, it also sets an individual latch associated with that error. These latches can be cleared only by powering down the drive or by pressing the maintenance FAULT-CLEAR switch on the fault card.

The following conditions can cause the fault latches to be set: write fault, head select fault, read and write fault, off-cylinder fault, and voltage fault.

Write Fault

A write fault is indicated by any of the following conditions:

- A low output from the write driver indicating that it might not be operating properly
- Low current input to the write driver
- Low voltage to the write driver
- No write data transactions when the write gate is active

Head Select Fault

A head select fault can be generated whenever more than one head is selected. If more than one head is selected, the circuit will generate a multiple select fault.

Read and Write Fault

A read and write fault is generated whenever the drive receives a READ gate and a WRITE gate simultaneously.

(Read or Write) Off-Cylinder Fault

A (read or write) off-cylinder fault is generated if the drive is in an off-cylinder condition and it receives a READ or WRITE gate from the controller.

Voltage Fault

A voltage fault is generated whenever the voltages are below a satisfactory operating level.

Errors Not Indicated by Fault Latch

The errors that are not indicated by a fault latch are detected by the disk storage device but are not stored in the fault latches. However, they do give other error indications.

A LOW SPEED or VOLTAGE FAULT signal occurs when the drive detects either a low voltage condition or that the drive spindle speed is below 75 percent of its necessary revolutions per minute. If either of these conditions are detected, the write drive circuits are disabled and a WRITE PROTECT signal is sent to the controller. This also results in an emergency retract of the heads.

No-servo track faults occur if dibits are not detected within 350 milliseconds after the load seek sequence begins. This lights a fault indicator and enables the return to zero logic. This also causes the heads to unload, and another load cannot be initiated until the no-servo tracks latch is cleared.

A seek error can occur under any of the following conditions:

- On-cylinder is not obtained within a prescribed time from the start of the seek
- Forward or reverse end of travel (EOT) has been sensed
- The drive is commanded to seek to a cylinder address greater than its boundaries

Setting this error latch enables the seek error line to the controller and inhibits the drive from performing another seek operation until the latch has been cleared. This latch can only be cleared by a return to zero seek command.

Block 12

Other Magnetic Storage Devices

Introduction to Mass Storage Devices

As its name implies, the mass storage subsystem (MSS) stores vast amounts of data. It is also capable of fairly fast access times. It is a magnetic tape system that utilizes a different type of tape. This activity introduces the mass storage subsystem (MSS).

Components of MSS

As can be seen in figure 12-1, the mass storage subsystem is a multicabinet system.

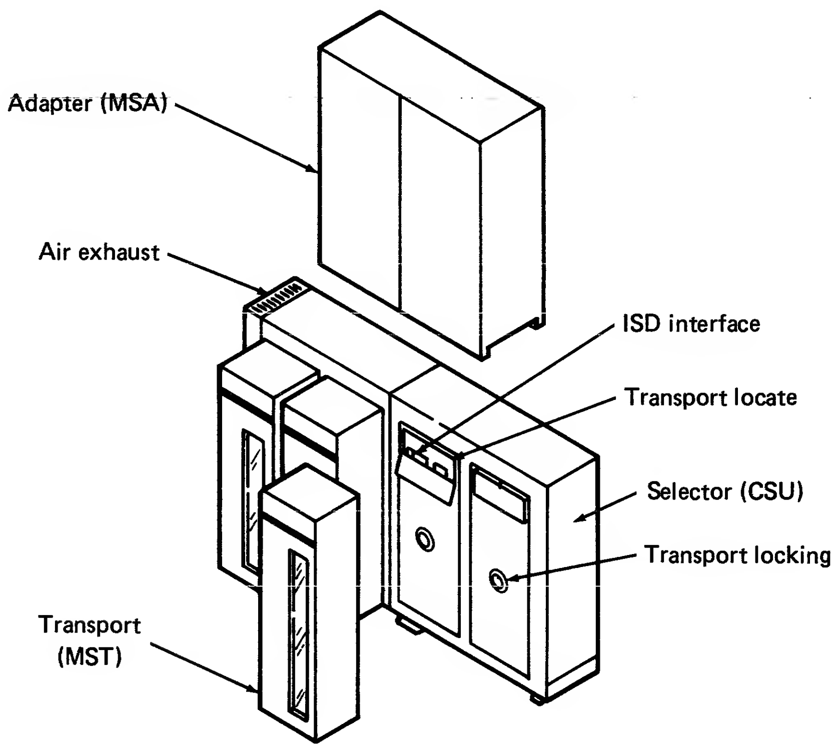

Figure 12-1. Mass storage subsystem

Tape cartridges are filled in the cartridge storage unit (CSU) according to an X-Y axis grid scheme. Any specific cartridge is selected by its mechanism and moved from storage to the entry station (input ISD) of the transport. Transport and CSU are under direct control from the mass storage adapter (MSA) which coordinates data, command, and status signals.

Once the transport detects a tape cartridge in its input ISD, status is sent to the adapter. The adapter then issues into the transport load station appropriate commands for cartridge motion. Load commands cause the transport to sequence the load procedure, whereby the cartridge door is opened, tape loads into the tape path and loop columns where it is positioned for the start of data transfer, and a magnetic head is indexed by an actuator. During the load operation, tape is removed from the cartridge in the form of a loop; i.e., both ends of the magnetic tape remain attached to the cartridge. One end of tape is secured to the cartridge spindle for loading and unloading and the other end is secured to the cartridge door.

During data transfer from or to the adapter, tape motion is controlled by a single capstan and loop positioning sensors located within the loop columns. Head and stream selection are directed by the adapter and functionally carried out by the head actuator servo system.

After data transfer, the adapter issues an unload command. The transport, acting upon this command, retracts its magnetic head, winds tape back into the cartridge, and closes the cartridge door. The cartridge is moved from the load station to an output ISD and is transported to an output station. The CSU then returns the cartridge to its addressed storage position. Operations of selection, transfer, reading, writing, and storage are overlapped to maintain a continuous flow of data between the transport and Central Processing Unit.

Transport Functional Groups

The transport includes the following functional groups as shown in figure 12-2:

- Control
- Cartridge handling
- Head actuator
- Tape drive
- Read/Write
- Pneumatics
- Power

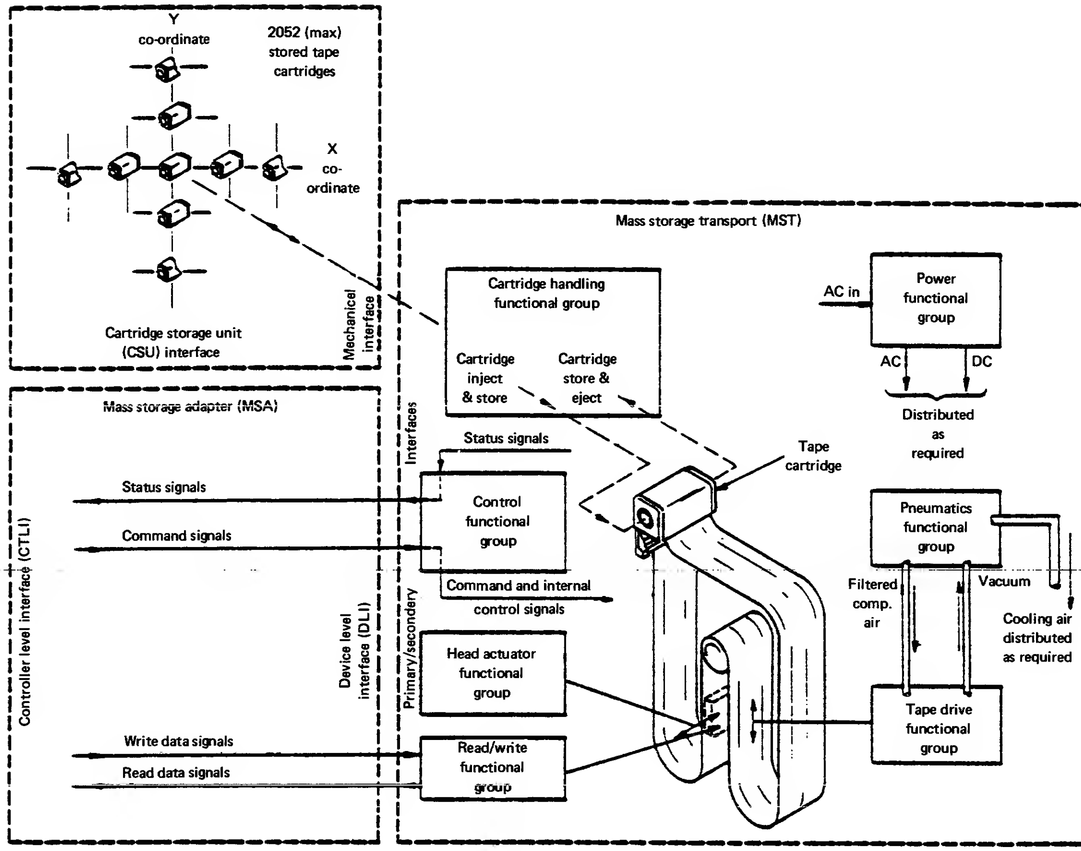

Figure 12-2. Mass storage transport orientation diagram

Figure 12-3 illustrates the complete transport in an overall block diagram form. Shown are the seven functional groups interrelated through the logic of the printed circuit boards of the control functional group. The drawing shows input/output signal interfaces and internal major signal flow.

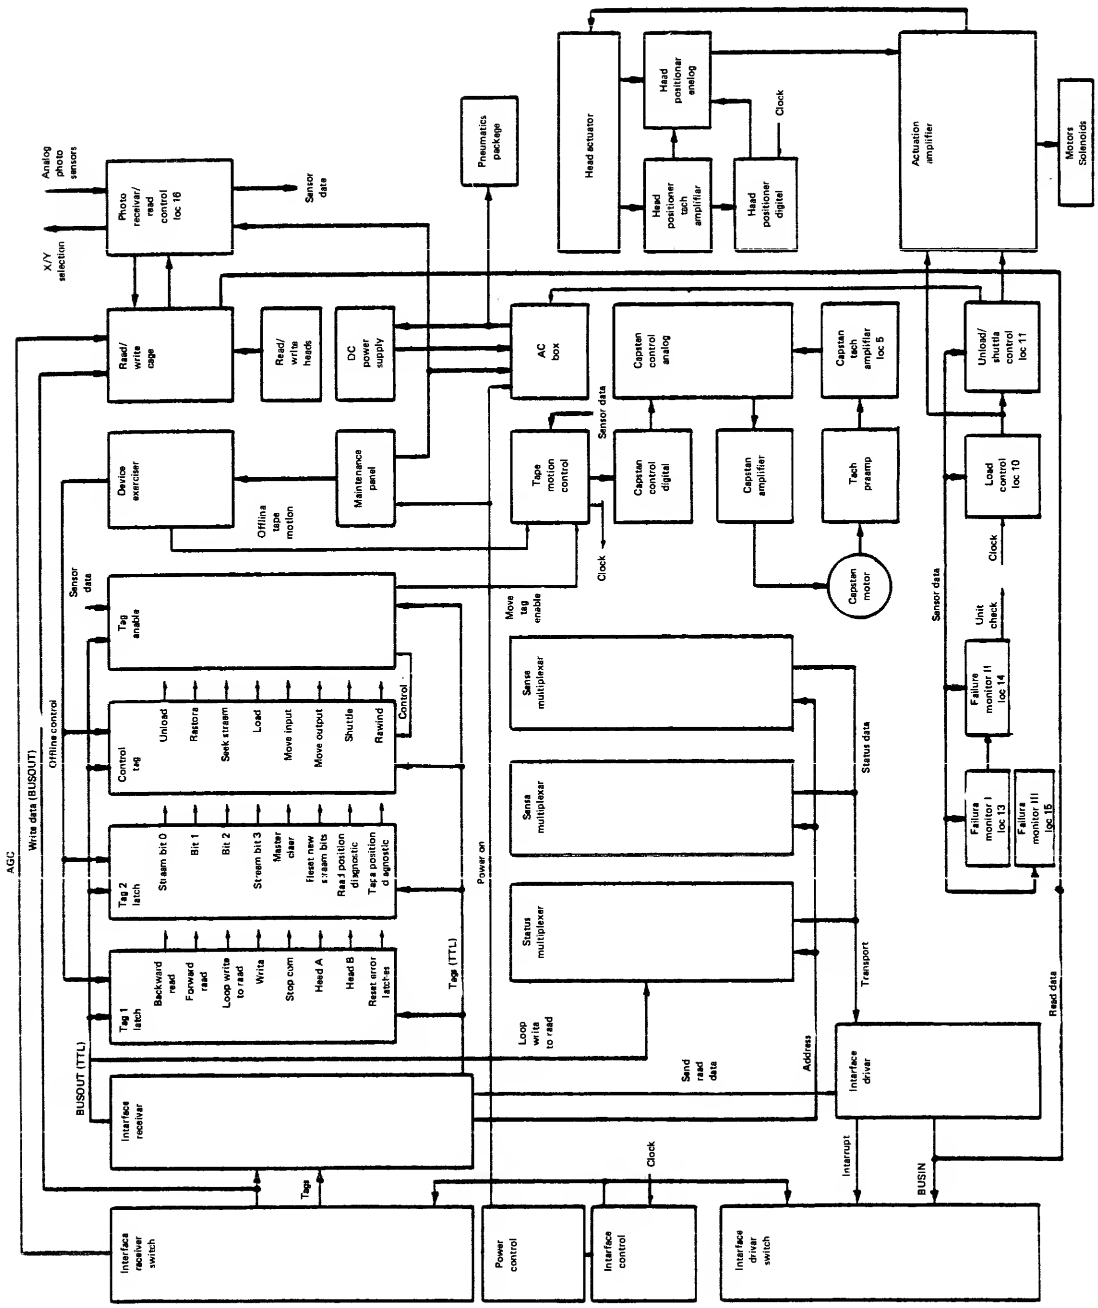

Figure 12-3. Transport overall block diagram

The control group directs the operation of the transport in response to external control commands and internal feedback. Control commands are received from the adapter, while feedback is received from the transport functional groups in the form of internal status signals.

The cartridge handling group moves tape cartridges from the selector to the tape path entrance for data processing, loads and unloads tape, and then transports the cartridge back to the CSU. Each tape transport is associated with an input and output intersta-tion drive. Five cartridges may be circulating at one time: one just being put into the entrance port (station A), one waiting in load and queue (station B), one in the trans-port being read and/or written on (station C), one in the exit storage awaiting return to the CSU (station D), and one at the exit port pick-point of the CSU (station E).

The head actuator group positions the read/write head at any desired stream location on the tape. The width of the tape and the large number of tracks require an indexable head. The head assembly employs two individual heads, one for the first 9 tracks (col-lectively called a stream) of data during the first tape traverse in one direction, and the other for the second 9-track stream of data during tape movement in the opposite direc-tion. The two sets of streams are called a stream pair. Since 144 tracks are recorded across 2.7 inches of tape width, tape movement is required eight times in both directions (with eight head-index movements) to read and/or write data on all 144 tracks. Write operations (with automatic read-after-write) are performed in one direction only, while read may be accomplished in either direction (backward read being employed for spac-ing over a preceding block of data).

The tape drive group moves the loaded tape past the read/write head in the tape path. The tape drive also accurately stops the tape in a controlled manner in the interblock gaps (IBG). A single capstan is used to move the tape in both directions across the magnetic head.

The read/write group transfers the data between the tape and external system via the adapter interface. This group consists of preamplifiers, read amplifiers, data receivers, write drivers, read/write head, and the necessary data-transfer control circuitry. There are no amplitude or skew adjustments within the transport itself. Skew and clip level requirements are controlled within the adapter.

The pneumatics group generates and distributes the various vacuum and air pressure requirements for the transport. The pneumatics package consists of a drive motor con-nected to a vacuum blower and air compressor, a plenum chamber, and cooling blower for other pneumatics package components. The compressor provides vacuum that: 1) draws tape from the opened tape cartridge into the vacuum columns, 2) maintains the tape loops in a uniform or slightly taut position within the loop columns, and 3) provides cooling air to the capstan motor. The air compressor provides filtered air to the air bearings. Ventilation exhaust for the cooling blower is provided at the top of the unit.

The power group accepts three-phase (60 or 50 Hz) electrical power, distributes the AC, and converts it to the various DC operating voltages required to power the electronics and drive motors.

Control Functional Group

The control functional group performs the following functions:

- Controls the operation of the transport in response to decoded commands from the adapter
- Monitors internal operation for status and faults, and initiates appropriate action
- Collects and encodes internal status conditions for transmission to the adapter
- Receives and transmits data for the media via the I/O circuits
- Provides on-line/off-line control and related operations for the transport

The operation of the transport is asynchronous, by which, broadly speaking, the completion of one operation initiates another. The command and status signals interact through the internal logic circuitry to: 1) produce commands necessary to control, initiate, or time the various operations, and 2) generate status signals signifying the state of these operations or related faults. The status signals are then encoded and transmitted to the adapter.

During off-line operation for checkout, maintenance, or training, the role of the adapter is partially simulated by a maintenance panel which is part of the control group. This panel includes switches for initiating certain control commands and capability for controlling application of power.

Cartridge Storage Unit

The cartridge storage unit (see figure 12-4), referred to hereafter as the CSU, is one component in the mass storage subsystem.

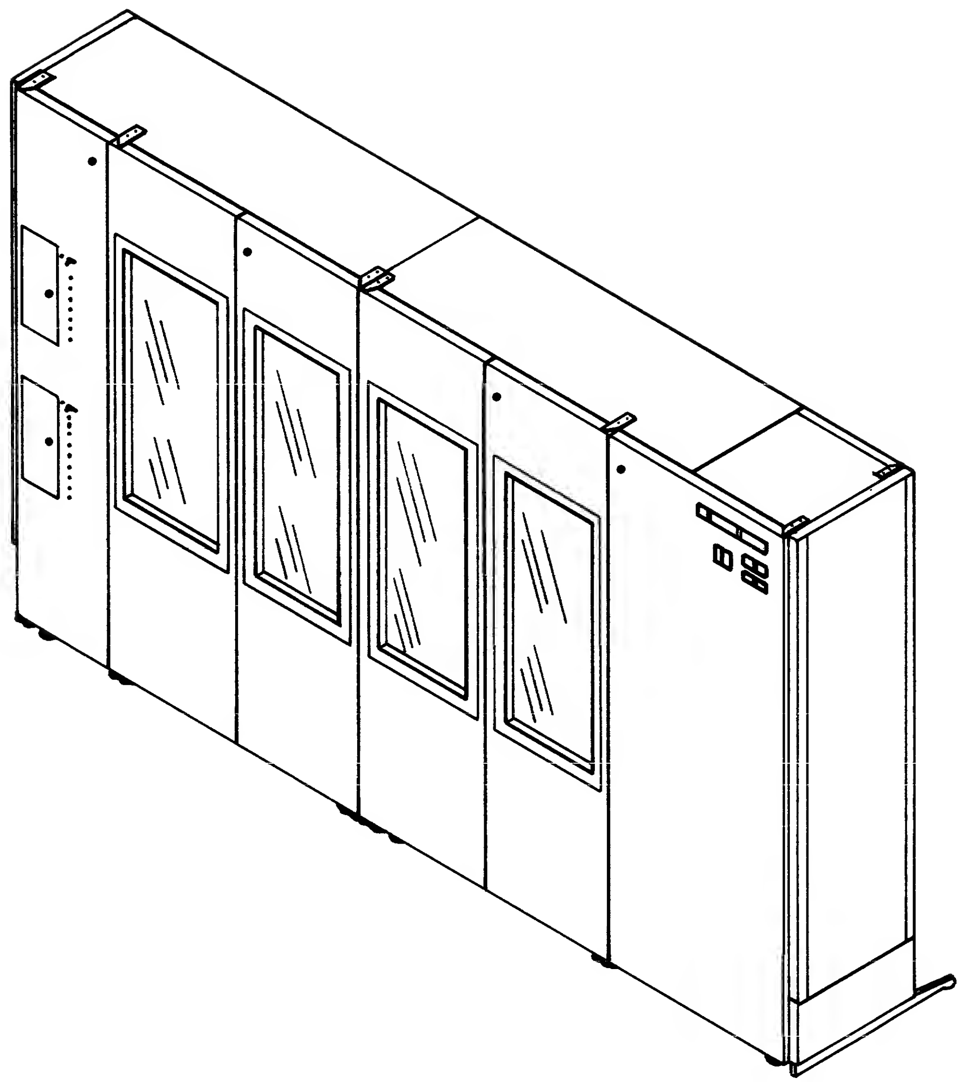

Figure 12-4. Cartridge storage unit

The function of the CSU within the subsystem is to store and deliver magnetic tape cartridges to one to four tape transports. It is physically and electronically connected to the mass storage transport (MST) and is controlled by a mass storage adapter (MSA). The CSU presents a status of its operational condition to the MSA upon request.

For purposes of physical orientation and common reference within this activity, the CSU and MST(s) are considered to be placed back to back. Thus, the front of the CSU pertains to the main door area through which the picker head and cubicles are accessed with left and right established from the position of viewer.

The unit contains stationary cubicles for the storage of 2052 tape cartridges plus eight cubicles on each of two input/output modules. Access to this storage matrix is normally accomplished by use of an input/output module to ensure security of the data base. Cartridges are moved within the CSU by a picker assembly mounted on a moving X/Y positioner. On command, the unit will retrieve a cartridge from the input/output module or storage matrix and move it to the inject interstation drive (ISD) of the selected MST. After processing by the MST, the cartridge is picked up at the eject ISD and moved back to the matrix for storage or to the input/output module for removal from the CSU. The maximum cycle time for any single operation is five seconds. This includes traversing between any two positions, retrieving the cartridge from or inserting the cartridge into the selected cubicle.

The CSU normally performs in an on-line mode of operation with an MSA and, if required for servicing, performs in a manual off-line mode through its maintenance panel.

The storage matrix, X/Y positioner, and cartridge picker mechanism are some of the major components in the CSU (see figure 12-5).

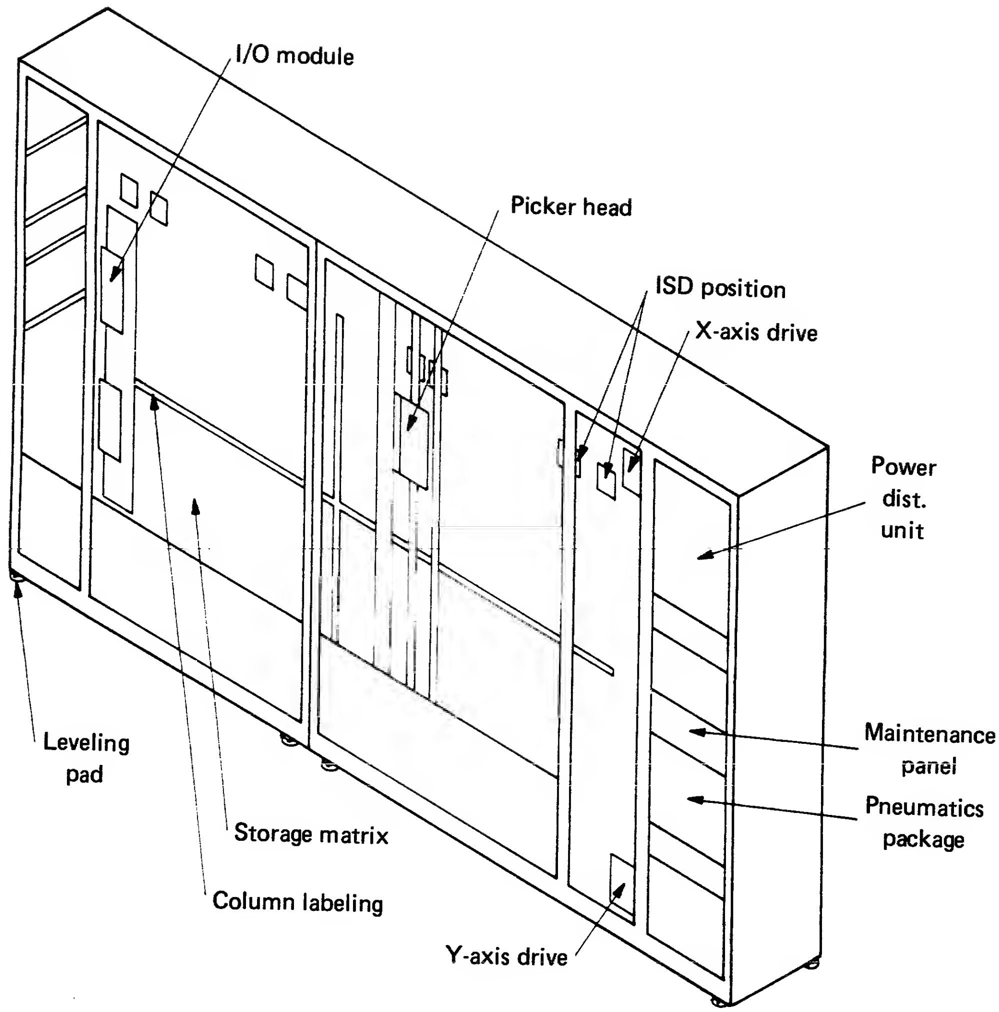

Figure 12-5. Major components of the cartridge storage unit

Storage Matrix

The storage matrix is a bank of 144 cubicle strips of 18 each (2052 cubicles total) used to file the cartridges stored in the CSU. The locations on the grid are numbered for easy identification. The cartridge storage cubicles are slanted at five degrees from the horizontal for gravity assist in retaining the cartridges.

X/Y Positioner

The X/Y positioner is an electromechanical device that carries the cartridge picker mechanism to the various locations on the storage matrix interstation drive ports and input/

output modules. A motor located to the upper right of the storage matrix drives the positioner along the X-axis (horizontal) through a system of timing belts and pulleys. Motion is transferred by a torque tube from the upper to the lower drive system. The positioner traverses on a upper guide rail and lower splined shaft and X-way that run along the base of the unit. The splined shaft is driven by a second motor located to the lower right of the storage matrix and it, in turn, drives an additional timing belt and pulleys to move the picker head along the Y-axis (vertical). A quadrature tachometer housed within each drive motor creates signals that are decoded to determine the X/Y position of the picker head.

The X/Y positioner provides the CSU with a means to transfer a cartridge within the physical limits of the storage module matrix, I/O module input station, and MST(s) pick-up and drop-off ports. The X/Y positioner consists of the following subassemblies: X- and Y-carriages, X- and Y-drive mechanisms, X- and Y-axis servo power amplifiers, X- and Y-axis position control servos, and servo power supply.

X-Carriage. The X-carriage subassembly (figure 12-6) contains the picker mechanism cabling, protective shroud, and components to support the Y-carriage. Within the picker mechanism cabling are power and signal cables and three air lines for pneumatic operation of the picker mechanism. Protection and guidance for the cabling during motion of the X-carriage is provided by the shroud. The upper and lower pulleys/bearings, Y-support tube, and Y-guide subassemblies of the X-carriage provide a means of support for rotational requirements of the Y-carriage. Tension of the Y-axis timing belt is afforded by the upper pulley/bearing subassembly. The lower pulley/bearing subassembly is a compound rotary, linear motion bearing that permits simultaneous X- and Y-axis motion. Stability, guidance, and position of the Y-carriage relative to the plane of the storage matrix is provided by the Y-support tube and Y-guide of the X-carriage subassembly.

Y-Carriage. Major components of the Y-carriage (figure 12-6) are picker mechanism (pick/put arm and cylinder subassemblies), air jet reservoir and solenoid, pick and put solenoids, position sensors, and casting and bearings. Picker mechanism retrieves cartridges from the cubicles or MST(s) upon activation of air jet solenoid, or drops cartridges off at selected destinations. Position sensors detect the following:

- Picker head back (PHB)
- Picker head full forward (PHFF)
- Put arm forward (PAF)
- Put arm back (PAB)
- Matrix probe (MPRB)
- Cartridge in picker head (CIPH)

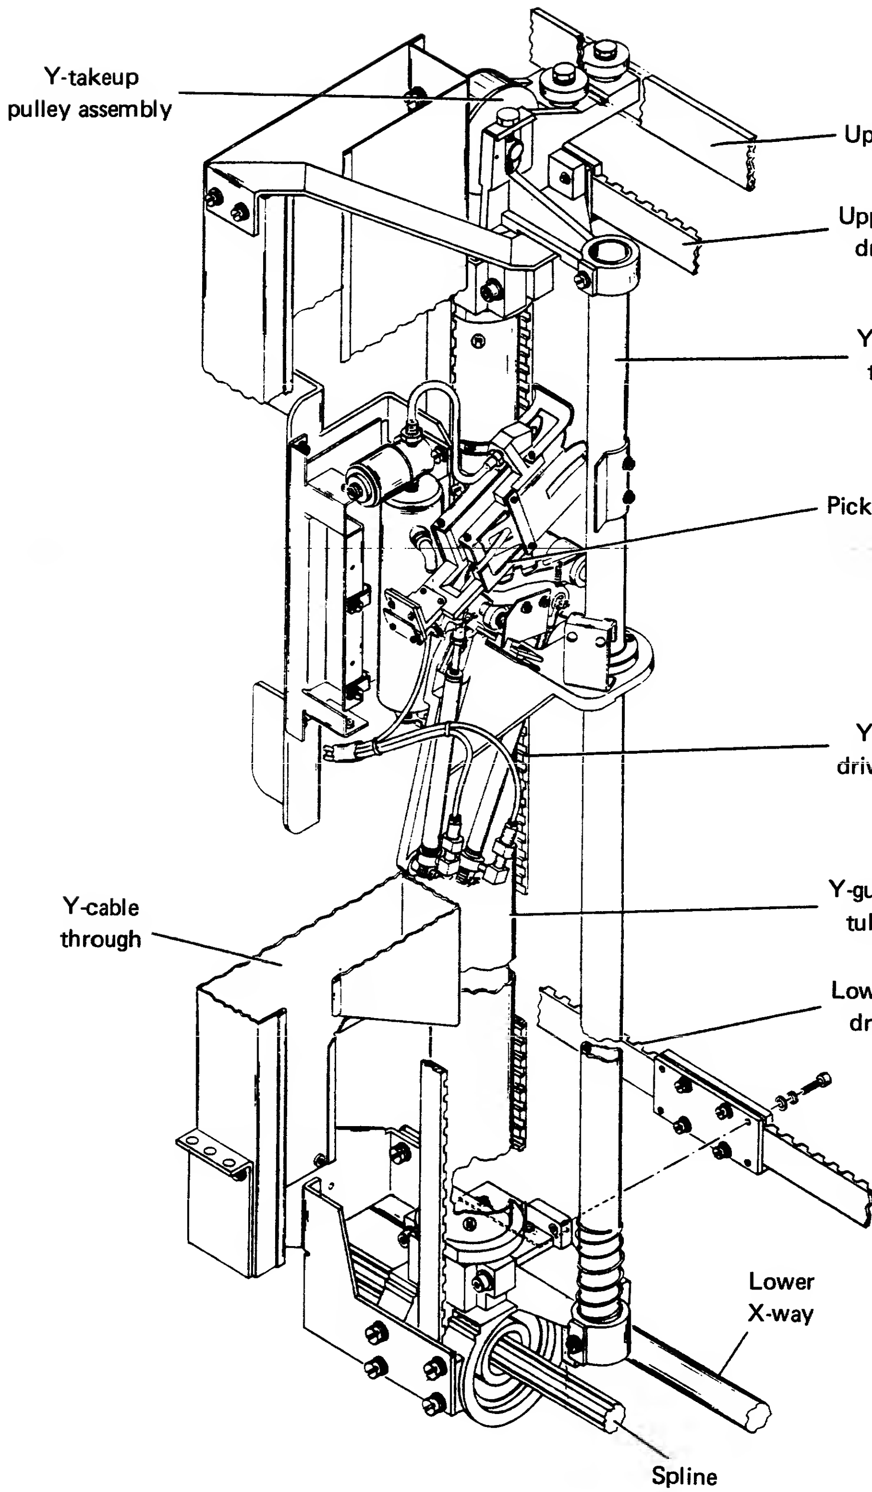

Figure 12-6. X/Y carriage assemblies

X-Drive Mechanism. The X-drive mechanism is powered by a DC servo-controlled motor having sufficient power to accelerate the inertia of the mechanism, overcome system friction, and attain traverse time consistent with the five-second pick/put cycle (see figure 12-7). A shaft encoder is an integral part of the motor assembly and contains a photodisk with 800 lines per revolution to provide the required linear distance resolution. In addition, the encoder supplies two square wave signals (one lags the other by 90 degrees) for determination of direction (left or right).

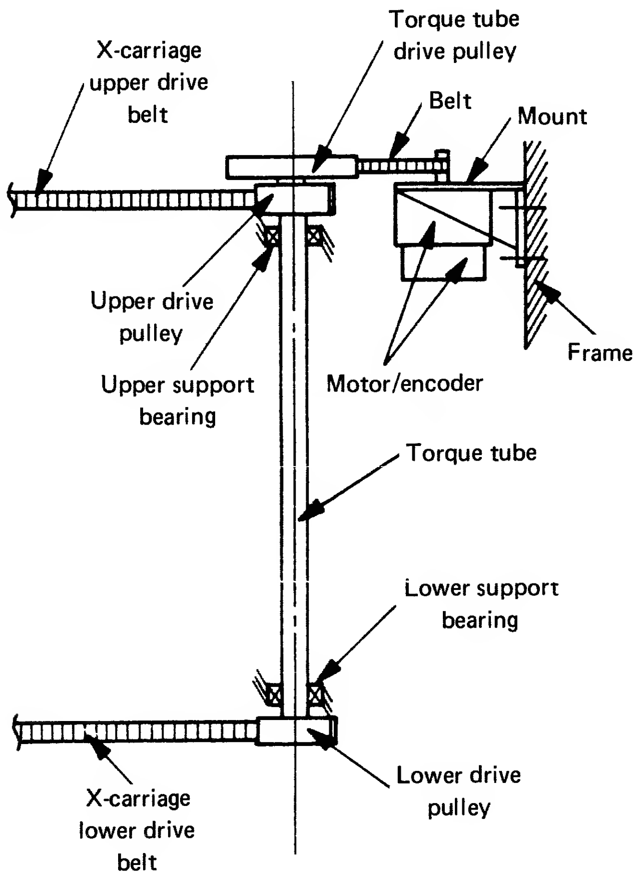

Figure 12-7. X-axis drive (simplified)

The spline and X-way located in the base of the CSU and the X-guide at the top of the CSU provide both support and drive linkage to the X-carriage. The spline allows motion of the X-carriage in the X-direction and imparts linear motion to the Y-carriage in the Y-direction. Double ball bearings at each end of the spline provide support and minimize whipping action when the X-carriage is traveling near either end. Lubrication of the spline is accomplished by a wick-wiping device.

A torque tube subassembly mounted vertically on the right side of the CSU transfers X-motion between the upper and lower X-drive belts. The two belts are the timing type and run in synchronism to eliminate skew with reference to the plane of the storage matrix. Top and bottom mountings of the torque tube allow for adjustment of the two X-axis drive belts relative to each other. Tension of both belts is adjustable, as well as the short belt between the torque tube assembly and X-motor.

Y-Drive Mechanism. The Y-drive mechanism is powered by a DC servo-controlled motor having sufficient power to accelerate the inertia of the mechanism, overcome system friction, and support the picker mechanism to attain a vertical traverse time of 1.6 seconds upward and 1.2 seconds downward. A shaft encoder is an integral part of the motor assembly, containing a photodisk with 800 lines per revolution to provide the required linear distance resolution. Two square wave signals from the encoder are displaced 90 degrees and determine the direction (up or down). The motor drives the spline shaft through a timing belt and the shaft transfers motion to the Y-carriage through a second pulley/belt arrangement. Alignment and belt tensioning are adjustable at the upper pulley mount.

X/Y Addresses. The position of the X/Y positioner is controlled by binary coordinates that originate in the MSA. Each position is identified in two parts: an X- and a Y-address. X-addresses are numbered consecutively from right to left, starting with X = 0. The maximum X-address is a binary 50 which corresponds to the "in" position of the I/O drawers. Y-addresses are numbered consecutively from bottom to top, starting with Y = 0.

Picker Mechanism (Head)

The picker mechanism (head) is a pneumatic electromechanical device that "picks" a cartridge from the selected location, holds it during transfer, and "puts" the cartridge into a new location. Mechanical linkages actuated by air cylinders with spring returns are used to position the picker receptacle for pick/put operations. A put arm pushes the cartridge out of the receptacle during a put and an air jet places the cartridge into the receptacle during a pick.

Air pressure is used to perform three functions in the CSU. Solenoid valves identified as picker head air solenoid valve (PHASV) and put arm air solenoid valve (PAASV) gate pressure to the cylinders on the picker head for actuation of the respective arms to accomplish pick/put operations. The valves are located on the picker head and air lines are routed to the picker head with the electrical cabling. The third valve, air jet air solenoid valve (AJASV), is mounted on the picker mechanism along with an air reservoir to help maintain the force of the air jet when performing a pick operation.

Introduction to the Drum

This activity introduces the magnetic drum storage device. Both the drum and devices are covered in this activity.

Magnetic Drums

The magnetic drum memory consists of a thin metallic cylinder made of iron or aluminum which is coated on the outside with a thin magnetic coating (see figure 12-8).

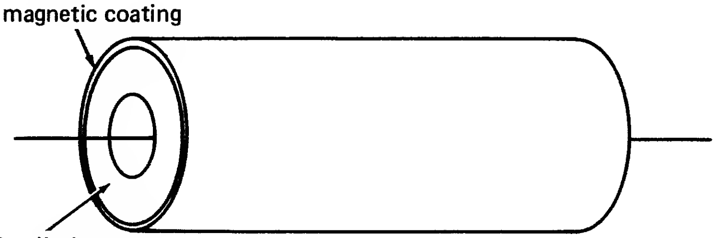

Figure 12-8. Magnetic drum construction

The drum memory stores the binary digits 1 and 0 in the form of magnetic spots on its thin magnetic coating. Drum memories are used instead of core memories when access time is of less importance than the storage capacity, and when cost is an important factor. In general, drum memories are considerably cheaper than their magnetic core counterparts; however, the drum memory is considerably slower. Typically, the drum memories in use today have an access time ranging from 10 milliseconds to as much as 1 second. Compare this to the access time of magnetic core memories which may be 1 microsecond or less. As you can see, magnetic core memory is significantly faster than magnetic drum memory. But also consider that magnetic core memory is 1000 times more expensive than magnetic drum memory. This is a classic example of the dilemmas associated with computers; speed of processing versus cost of equipment.

Magnetic Drum Devices

In the drum memory, a magnetic recording head, similar to those used on home tape recorders, is used to read and write data onto the magnetic film coating on the drum (see figure 12-9). As the drum surface passes the read/write head, a current passing through the writing head causes a spot on the drum surface directly beneath the head to be magnetized.

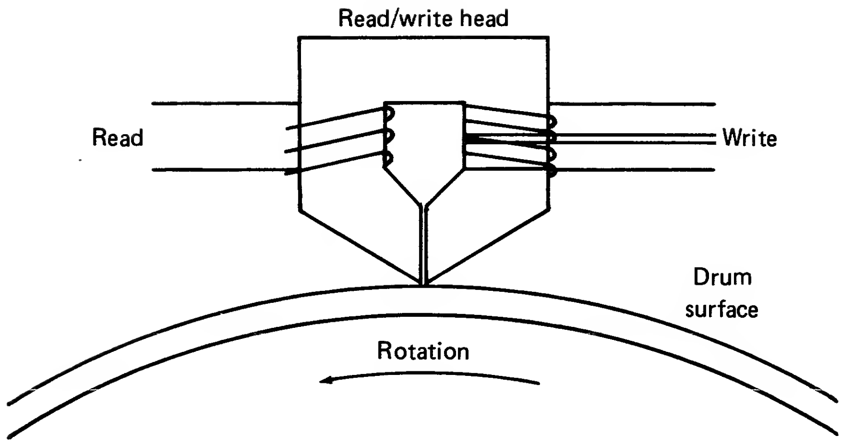

Figure 12-9. Read/write head

Similarly, when reading data out of the drum memory, the magnetized spots on the drum surface cause a current to be induced into the reading head when the spots on the drum surface are directly below the head.

When writing the data bit on the drum surface, a certain small area of the drum is magnetized. This area is called a cell and is the area required to store a single bit of data (see figure 12-10). Obviously, the smaller the cell area required to store a bit of data, the larger the number of bits that can be stored on a given drum surface. This again is a question of economics; the smaller the cell area, the more costly the read/write heads and associated equipment.

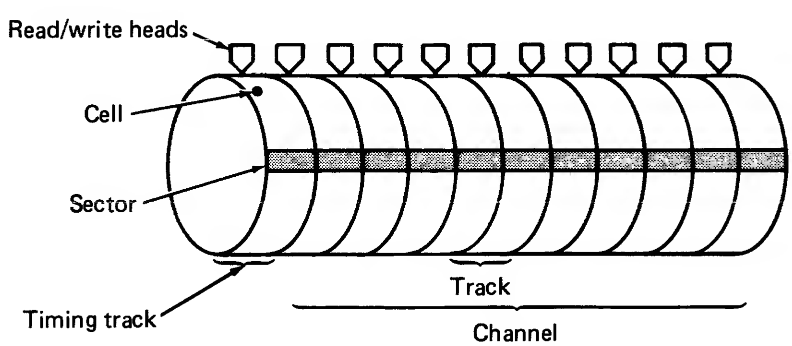

Figure 12-10. Drum memory format

One drum revolution, therefore, contains a certain given number of cells. This strip of cells, which rings the circumference of the drum, is called a track (see figure 12-10). The line of cells that is parallel to the center of the drum is defined as a sector. A computer word is usually recorded in one of these sectors so that the number of cells in a sector is the computer word length. The number of cells in a track, therefore, defines the word capacity of the drum. For example, in figure 12-10, there are ten cells in the sector. (The end track is usually a timing track that has clock pulses recorded on it for purposes of circuit timing.) In addition, assume that there are 16 cells around the outside of the drum in one track. The drum, therefore, is capable of storing 16, 10-bit words.

In the core memory, access to a given word requires only addressing the particular word and the data can be read out or written in almost instantly. In the case of the drum, however, after addressing the sector, the read/write heads are not activated until the proper sector appears under the heads. Since the drum rotates at a constant speed, the amount of time for a particular sector to appear under the heads can vary. That is, if the next sector to pass under the heads is the addressed sector, the access time is relatively short. On the other hand, if the addressed sector is the sector that has just passed by the heads, the heads do not become active until the drum makes almost one more complete revolution. In the second case, the access time is relatively long. Since each condition is equally probable, the access time of the drum is specified as the time that the drum will take to make one-half revolution. For example, if a drum rotates once every 10 milliseconds, the access time would be 5 milliseconds.

In the example shown in figure 12-10, the computer word filled up the entire length of the drum. To obtain additional storage capacity, the drum can be extended along its axis to accommodate more computer words. In the previous example, the drum was capable of holding 16, 10-bit words. By doubling the length of the drum, 32, 10-bit words can be handled. This is done by defining the area of the basic drum as one channel and the additional area as the second channel. Addressing a particular word in memory proceeds as it did for the previous example, but now the proper channel must be included as part of the addressing information.

Figure 12-11 shows a four-channel system. Note that the timing track is not repeated in each channel. In this case, four times as many read/write heads are also required, which can get rather expensive. To alleviate this expense, modern systems use only the original set of read/write heads and merely move the heads over the addressed channel. Although this allows a considerable cost saving, access time now becomes greatly increased because additional time is required to position the read/write heads. Also, some sort of positioning device is required. All of these factors must be taken into consideration and weighed when a drum memory is selected.

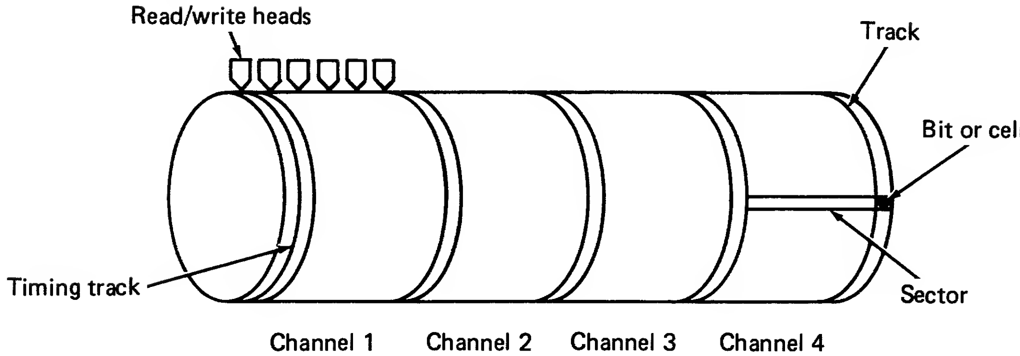

Figure 12-11. Four-channel drum memory

PROGRESS CHECK

1. What happens to a piece of wire when a current passes through it?

    a. It becomes a permanent magnet.
    b. It becomes a temporary magnet.
    c. A small magnetic field builds up around it.
    d. It overheats and melts.

2. The read head used in magnetic surface recording uses the principle of __________.

    a. flux density
    b. high permeability
    c. low voltage
    d. induction

3. Phase modulation records a flux change __________.

    a. only on zero bits
    b. only on one bits
    c. that is negative on zeros and positive on ones
    d. only on one bits and between each bit cell

4. A binary code consisting of seven bits can represent __________ different codes.

    a. 64
    b. 128
    c. 192
    d. 256

5. A __________ is the data recorded on tape by a single recording head.

    a. track
    b. file
    c. record
    d. frame

6. On a magnetic tape drive system, which error detection method counts all the bits in each track?

    a. CRC
    b. LPC
    c. LRC
    d. BCD

PROGRESS CHECK

7. Which of the following statements is true?

   a. All transports have at least two capstans.
   b. Tape transports primarily apply vacuum for cleaning tape.
   c. All transports supply tape from the right and take up tape on the left.
   d. Some tape capstan motors will operate in both forward and reverse directions.

8. What is the third step to perform while manually loading a tape transport?

   a. Manually thread tape onto take-up reel.
   b. Close window and press load button.
   c. Open window and mount tape.
   d. Move tape to BOT marker.

9. Which of the following statements is true?

   a. The movement of tape across the head should be independent of reel control.
   b. The movement of tape across the lead should be a function of reel speed.
   c. Usually audio tape motion is more precise than digital tape motion.
   d. Digital tape motion is slower than audio to increase the bits/inch density.

10. Which type of braking system is simplest, fastest, and easiest to maintain?

    a. Mechanical
    b. Pinch roller
    c. Vacuum brake
    d. Pressure pad

11. Optical tachometers __________.

    a. generate voltage proportional to tape speed
    b. are used only during high speed rewinds
    c. generate pulse rates proportional to tape speed
    d. use holes in tape reels for timing marks

PROGRESS CHECK

12. Photo signal oscillation or excessive triggering may be caused by __________.

    a. a poorly aligned light source
    b. low amplifier gain
    c. excessive amplifier gain
    d. a weak light source

13. Which of the following statements is true?

    a. Write current flows at all times with one polarity with RZ recording techniques.
    b. Write current flows at all times with one polarity with NRZ recording techniques.
    c. Write current flows at all times with reversing polarity with NRZ recording techniques.
    d. Write current flows at all times with reversing polarity with RZ recording techniques.

14. How many BOT markers are used on a 2400-foot reel of tape?

    a. One
    b. Two
    c. Three
    d. Four

15. Organize the data groups in ascending order, from small to large.

    a. Bytes, tracks, sectors, cylinders
    b. Bytes, sectors, tracks, cylinders
    c. Bytes, cylinders, tracks, sectors
    d. Bytes, cylinders, sectors, tracks

16. To what does capacity of a disk storage device refer?

    a. Speed of head actuator
    b. Number of heads on device
    c. Number of bytes device can store
    d. Thickness of iron oxide coating

PROGRESS CHECK

17. Which recording technique is described by the following statement? A one is recorded by a positive current change. A zero is recorded by a negative current change. This necessitates a negative current change between consecutive one cells, and a positive current change between consecutive zero cells.

    a. FM
    b. MFM
    c. PM
    d. DPRZ

18. What is the area of the disk surface that resembles a piece of pie?

    a. Track
    b. Cylinder
    c. Sector
    d. Spindle

19. On a disk storage device, which disk actuator moves the carriage in and out using a rack and pinion arrangement?

    a. Hydraulic
    b. Printed circuit motor
    c. Linear voice coil
    d. Rotary motion voice coil
    e. Stepper motor

20. On a disk storage unit using a stepper motor, when the stepper motor receives a pulse, how many tracks will it move?

    a. 1
    b. 5
    c. 20
    d. 100

21. Which one of the following statements is true?

    a. Bus-in consists of one signal line that is sent to each disk device controlled by the same controller.
    b. Tag lines include cylinder select tag, control select tag, and servocontrol tag.
    c. The unit select bus consists of multiple control lines.
    d. Tag line signals originate from the disk storage device.

PROGRESS CHECK

22. The two inputs to the subtractor in the servosystem of a disk drive are from the __________.

    a. cylinder address register: difference counter
    b. present address register: D to A converter
    c. cylinder address register: present address register
    d. D to A converter: demodulator

23. In a mass storage system using tape cartridges, how many cartridges can be in circulation at any given time?

    a. One
    b. Two
    c. Five
    d. No limit

24. In a mass storage system, the storage matrix bank of the CSU contains __________ cubicle strips.

    a. 2052
    b. 144
    c. 18
    d. 1

PROGRESS CHECK ANSWERS

1. Correct Answer: c
   Resource: Text *Magnetic storage Concepts*, Page 1-3

2. Correct Answer: d
   Resource: Text *Magnetic storage Concepts*, Page 1-3

3. Correct Answer: c
   Resource: Text *Magnetic storage Concepts*, Page 2-4

4. Correct Answer: b
   Resource: Text *Magnetic storage Concepts*, Page 2-12

5. Correct Answer: a
   Resource: CBE "Tape Format Introduction"

6. Correct Answer: c
   Resource: Text *Magnetic storage Concepts*, Page 3-8

7. Correct Answer: d
   Resource: Text *Magnetic storage Concepts*, Page 4-1

8. Correct Answer: a
   Resource: Text *Magnetic storage Concepts*, Page 4-2

9. Correct Answer: a
   Resource: Text *Magnetic storage Concepts*, Page 5-1

10. Correct Answer: c
    Resource: Text *Magnetic storage Concepts*, Page 5-1

11. Correct Answer: c
    Resource: Text *Magnetic storage Concepts*, Page 6-8,6-9

PROGRESS CHECK ANSWERS

12. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 6-18

13. Correct Answer: c
    Resource: Text _Fundementals Of Digital Magnetic Tape Units_, page 6-4,6-5

14. Correct Answer: a
    Resource: Text _Magnetic storage Concepts_, page 7-4

15. Correct Answer: b
    Resource: Text _Magnetic storage Concepts_, page 8-1

16. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 8-7

17. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 9-25

18. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 9-42

19. Correct Answer: a
    Resource: Text _Magnetic storage Concepts_, page 10-2

20. Correct Answer: a
    Resource: Text _Magnetic storage Concepts_, page 10-7

21. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 11-3

22. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 11-13

23. Correct Answer: c
    Resource: Text _Magnetic storage Concepts_, page 12-5

24. Correct Answer: b
    Resource: Text _Magnetic storage Concepts_, page 12-9

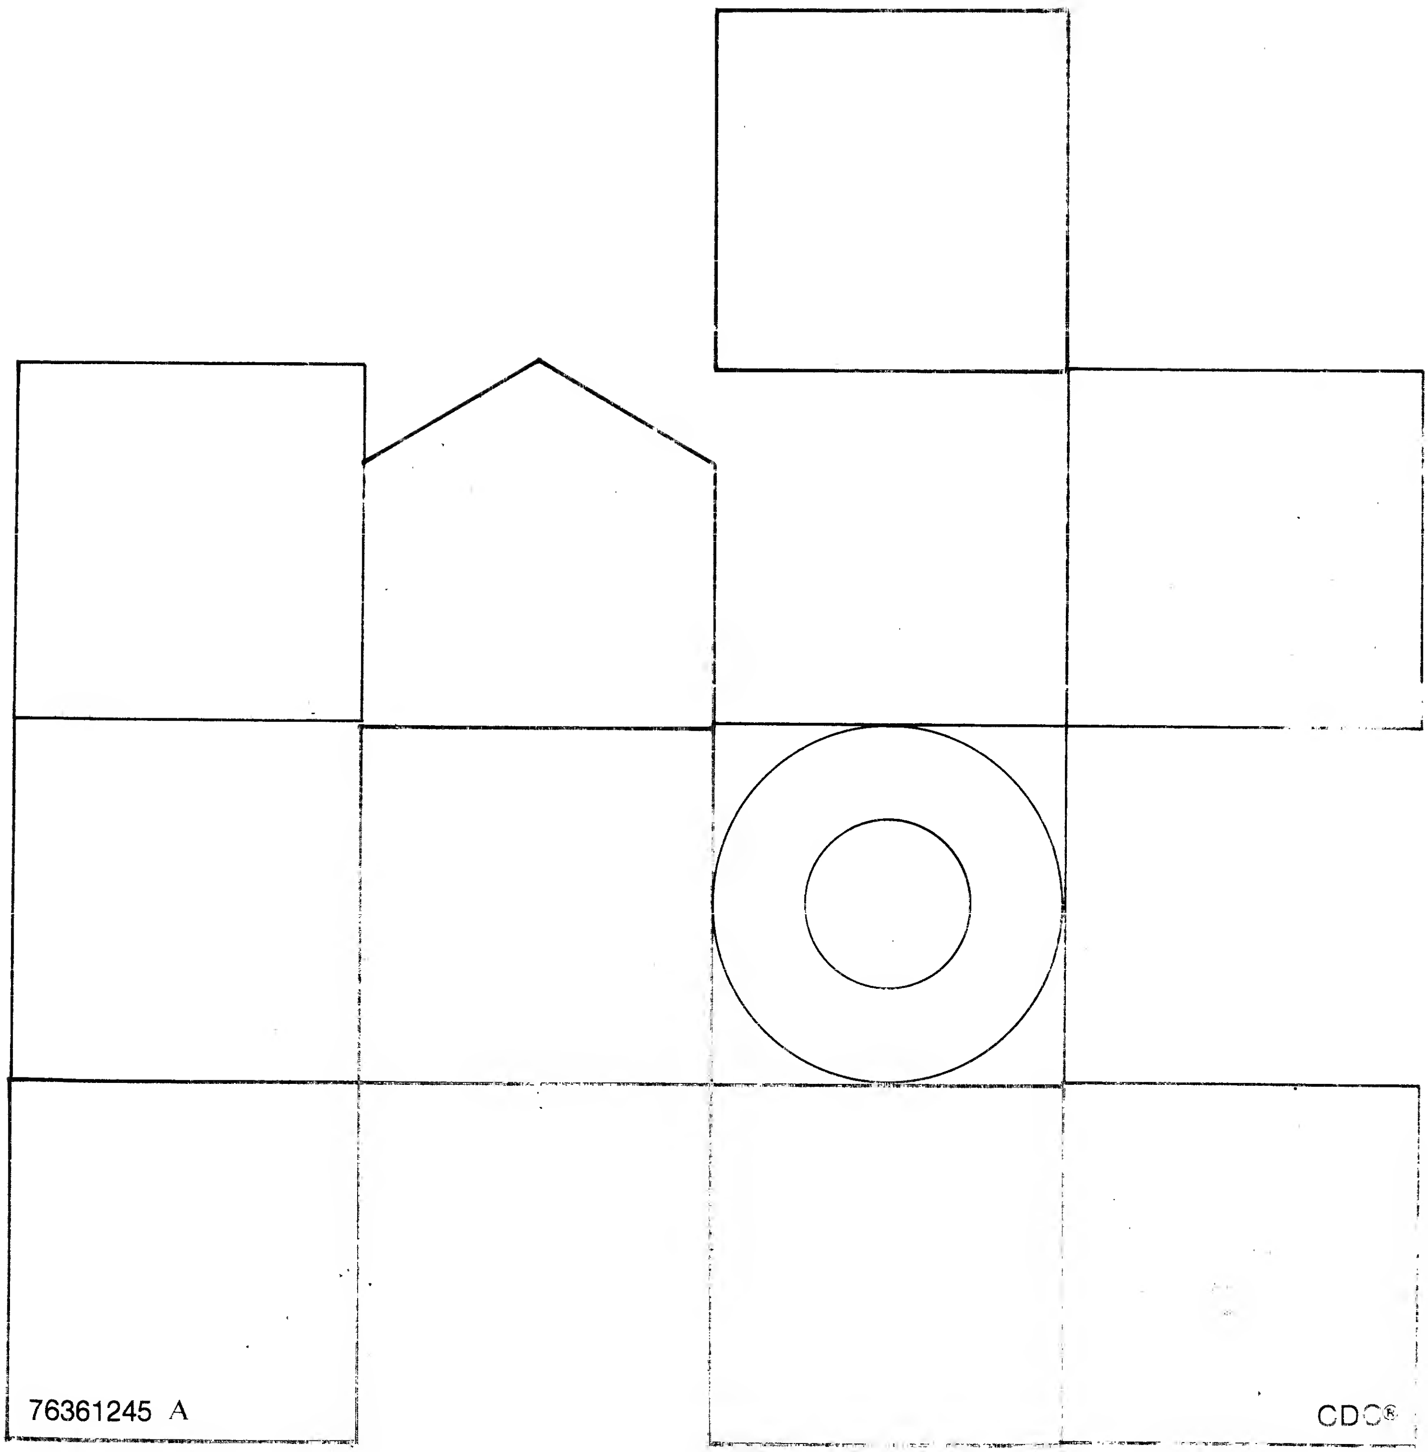
76361245 A
CDC®